

Vapnik-Chervonenkis Density in some Theories without the Independence Property, II

Matthias Aschenbrenner, Alf Dolich, Deirdre Haskell, Dugald Macpherson, and Sergei Starchenko

For Anand Pillay, on his 60th birthday.

Abstract We study the Vapnik-Chervonenkis (VC) density of definable families in certain stable first-order theories. In particular we obtain uniform bounds on VC density of definable families in finite U-rank theories without the finite cover property, and we characterize those abelian groups for which there exist uniform bounds on the VC density of definable families.

CONTENTS

1. Introduction	2
2. Preliminaries	5
3. Linear VC Density for Finite Rank Theories	9
4. Modules of Finite Breadth	16
5. Abelian Groups with Uniformly Bounded VC Density	30
References	50
Acknowledgments	52

Printed April 4, 2012

2000 Mathematics Subject Classification: Primary, 03C45, 52C45; Secondary, 03C60, 03G10

Keywords: theories without the independence property; Vapnik-Chervonenkis density; finite cover property; modules; abelian groups

1 Introduction

The Vapnik-Chervonenkis (VC) density and its cousin, the VC dimension, are numerical parameters associated to any set system (i.e., a family of subsets of a given base set). VC classes, that is, set systems whose VC dimension is finite, have been investigated since the early seventies (beginning with [35; 37; 42]). Since then, the concept of VC dimension has found numerous applications in statistics, combinatorics, learning theory, and computational geometry. In model theory, the connection between VC classes and the theories without the independence property (NIP theories) introduced by Shelah [36] was first made explicit by Laskowski [24], who showed that a complete first-order theory is NIP if and only if every formula defines, in each model, a VC class. An important class of NIP theories are the ones which are stable. The class of NIP theories, however, also encompasses, among other examples, all weakly o-minimal theories and many interesting theories of expansions of valued fields. In recent years, not least due to the efforts of Anand Pillay (see, e.g., [18; 20; 21]), there has been some progress in extending the reach of the highly developed methods of stability theory into this wider realm of NIP theories.

In this paper we establish uniform bounds, in terms of the number of parameter variables, on the VC density of formulas in certain stable theories, continuing our study from [1], where we mainly focussed on unstable examples. Such uniform bounds may be interesting since the VC density is often an important parameter for the complexity of a family of sets, and is related to discrepancy and the size of ε -approximations [27], and to entropic dimension [2]. Moreover, uniform bounds on VC density are intimately connected with a strengthening of the NIP concept, called dp-minimality, which has recently received attention through the work of Shelah [39], Onshuus-Usvyatsov [28], Dolich-Goodrick-Lippel [9], and others. Before we state some of the main results of this paper, we briefly recall the relevant terminology. In the rest of this introduction, \mathcal{L} is a first-order language, T is a complete \mathcal{L} -theory without finite models, and \mathbf{M} is a model of T (with universe M). Given a tuple $x = (x_1, \dots, x_m)$ of pairwise distinct variables we denote by $|x| := m$ the length of x .

1.1 Shatter function, VC dimension, and VC density. Let $\varphi(x; y)$ be a *partitioned \mathcal{L} -formula*, that is, an \mathcal{L} -formula with the free variables of φ contained among the entries of the tuples $x = (x_1, \dots, x_m)$ and $y = (y_1, \dots, y_n)$ of pairwise distinct variables (also assumed to be disjoint). We call the x_i the *object variables* and the y_j the *parameter variables* of $\varphi(x; y)$. We obtain a set system

$$\mathcal{S}_\varphi = \{\varphi(M^m; b) : b \in M^n\}$$

on M^m , where $\varphi(M^m; b)$ denotes the subset of M^m defined in \mathbf{M} by the formula $\varphi(x; b)$. Given $A' \subseteq A \subseteq M^m$, we say that A' is *cut out from* A by φ if $A' = S \cap A$ for some $S \in \mathcal{S}_\varphi$, and we say that A is *shattered* by φ if every subset of A is cut out from A by φ .

For each non-negative integer t let $\pi_\varphi(t)$ be the maximal number of subsets cut out from a t -element subset of M^m . (So $\pi_\varphi(t) = 2^t$ if and only if there is a t -element subset of M^m shattered by φ .) The function $t \mapsto \pi_\varphi(t)$ is called the *shatter function* of φ ; it is routine to verify (see, e.g., [1, Section 3]) that π_φ is an invariant of the elementary theory T of \mathbf{M} . The *VC dimension* of φ (in T)

is the largest $d = \text{VC}(\varphi)$ (if it exists) such that $\pi_\varphi(d) = 2^d$; if there is no such d (i.e., if $\pi_\varphi(t) = 2^t$ for each t), then we also set $\text{VC}(\varphi) = \infty$. (We suppress the dependence on T in the notation for the shatter function of φ and concepts, like the VC dimension of φ , derived from it.)

Now if φ has finite VC dimension d , then by a fundamental combinatorial fact (proved independently in [35], [37] and [42]), the shatter function of φ is bounded above by a polynomial in t of degree d ; in fact, $\pi_\varphi(t)$ is bounded by $\binom{t}{\leq d} := \binom{t}{0} + \binom{t}{1} + \cdots + \binom{t}{d}$ (the number of subsets of $[t] := \{1, \dots, t\}$ having at most d elements). One defines the *VC density* $\text{vc}(\varphi)$ (in T) of a formula φ with finite VC dimension in T as the infimum of all real numbers $r \geq 0$ such that $\pi_\varphi(t)/t^r$ is bounded for all positive t ; if $\text{VC}(\varphi) = \infty$, then we put $\text{vc}(\varphi) = \infty$.

1.2 Independence dimension and VC duality. The partitioned \mathcal{L} -formula $\varphi(x; y)$ is said to have the *independence property* for \mathbf{M} if for every $t \in \mathbb{N}$ there are $b_1, \dots, b_t \in M^n$ such that for every $S \subseteq [t]$ there is $a_S \in M^m$ such that for all $i \in [t]$, $\mathbf{M} \models \varphi(a_S; b_i) \iff i \in S$. The structure \mathbf{M} is said to have the independence property if some \mathcal{L} -formula has the independence property for \mathbf{M} , and to be *NIP* (or to be *dependent*) otherwise. Clearly the independence property for φ only depends on $T = \text{Th}(\mathbf{M})$, so it makes sense to speak of a partitioned \mathcal{L} -formula having the independence property for the complete \mathcal{L} -theory T , and of T being NIP. The *dual* of the partitioned \mathcal{L} -formula $\varphi(x; y)$ is the partitioned \mathcal{L} -formula $\varphi^*(y; x)$ which is syntactically the same formula as φ , except that the roles of object and parameter variables are interchanged. Thus φ has the independence property for T if and only if φ^* has finite VC dimension in T . (It is also true, by another fundamental observation about set systems, that φ has the independence property for T if and only if φ has finite VC dimension in T ; cf., e.g., [1, Section 3].)

Given $B \subseteq M^n$, a complete $\varphi(x; B)$ -type is a maximal consistent subset of

$$\{\varphi(x; b) : b \in B\} \cup \{\neg\varphi(x; b) : b \in B\}.$$

We denote the set of complete $\varphi(x; B)$ -types by $S^\varphi(B)$. The shatter function $\pi_\varphi^* := \pi_{\varphi^*}$ of the dual formula of φ counts the number of complete $\varphi(x; B)$ -types over finite parameter sets B : the map

$$B' \mapsto \{\varphi(x; b) : b \in B'\} \cup \{\neg\varphi(x; b) : b \in B \setminus B'\}$$

defines a bijection from the collection of subsets of B cut out by φ^* onto $S^\varphi(B)$. Hence

$$\pi_\varphi^*(t) = \max \{|S^\varphi(B)| : B \subseteq M^n, |B| = t\} \quad \text{for each } t.$$

This interpretation of the shatter function allows us to transform the problem of estimating $\text{vc}(\varphi^*)$, which is one of bounding the asymptotic growth of π_{φ^*} , into the problem of counting the number of $\varphi(x; B)$ -types over finite parameter sets B . For the latter task, ready-made model-theoretic tools (like the local ranks employed in Section 3 below) are available. As already in [1], we also use a variant of the familiar notion of definable type; we combine it with additional combinatorial techniques, centered around the notion of breadth of a set system, and apply it, in this paper, to modules. (See Sections 4–5.)

1.3 Finite-rank theories. Our first main theorem gives a uniform bound on VC density of formulas for superstable theories of finite rank which do not have the finite cover property. From [38, Definition II.4.1] recall that T is said to have the *finite cover property* if there is an \mathcal{L} -formula $\varphi(x; y)$ with the finite cover property in T , i.e., so that for arbitrarily large $t \in \mathbb{N}$ there are $b_1, \dots, b_t \in M^n$ so that $\{\varphi(x; b_i) : i \in [t]\}$ is inconsistent but $\{\varphi(x; b_i) : i \in w\}$ is consistent for any proper subset w of $[t]$. If T has the finite cover property, then there is actually a formula $\varphi(x; y)$ with the finite cover property in T where $|x| = 1$ [38, Theorem II.4.4]. Every theory which does not have the finite cover property is stable [38, Theorem II.4.2]. In Section 3 we show:

Theorem 1.1. *Suppose T does not have the finite cover property and T has finite U-rank d . Then every partitioned \mathcal{L} -formula $\varphi(x; y)$ has VC density at most $d|y|$ in T (in fact, $\pi_\varphi(t) = O(t^{d|y|})$).*

So e.g., for p a prime or 0, every formula $\varphi(x; y)$ in the language of rings has VC density at most $|y|$ in the theory of algebraically closed fields of characteristic p . Other examples of theories to which Theorem 1.1 applies (including all expansions of groups having finite Morley rank) are provided in Section 3.4. The proof of this theorem, given in Section 3, uses the local ranks $R^m(-, \Delta, \aleph_0)$ of Shelah [38, Chapter II], here denoted by $R_\Delta(-)$. The absence of the finite cover property enters the picture as a necessary condition for the definability of the ranks R_Δ ; see Fact 3.3 below.

1.4 Abelian groups. We have an improvement over Theorem 1.1 for \aleph_0 -categorical abelian groups. For an abelian group A , always written additively, let

$$U(p, i; A) := |(p^i A)[p]/(p^{i+1} A)[p]| \quad (p \text{ prime})$$

be its Ulm invariants; here $(p^i A)[p]$ denotes the subgroup

$$(p^i A)[p] = \{a \in A : pa = 0 \text{ and } p^i b = a \text{ for some } b \in A\}$$

of A . For each prime p let $U_{\geq \aleph_0}(p; A)$ denote the set of $i \geq 0$ such that $U(p, i; A)$ is infinite. If A has finite exponent, then each $U_{\geq \aleph_0}(p; A)$ is finite, with $U_{\geq \aleph_0}(p; A) = \emptyset$ for all but finitely many p . With this notation we have:

Theorem 1.2. *Let A be an infinite abelian group of finite exponent, construed as a first-order structure in the language $\mathcal{L} = \{0, +\}$ as usual. Let $n_p := |U_{\geq \aleph_0}(p; A)|$ and $m_p := \max U_{\geq \aleph_0}(p; A)$ if $U_{\geq \aleph_0}(p; A) \neq \emptyset$ and $m_p := 0$ otherwise. Set*

$$d := \sum_p \min \{n_p, \lfloor m_p/2 \rfloor + 1\}.$$

Then every \mathcal{L} -formula $\varphi(x; y)$ has VC density at most $d|y|$ in A .

Since, in the context of the previous theorem, the Morley rank of the abelian group A is $\text{MR}(A) = \sum_p (n_p + \sum_{i \in U_{\geq \aleph_0}(p; A)} i) \geq d$, the bound obtained here is more precise than the one in the general Theorem 1.1. Moreover, the bound in Theorem 1.2 can be refined even further; since the resulting bound is somewhat technical to describe, we refer to the proof of this theorem in Section 4 below for the improved result. (However, even the crude bound stated in the theorem is optimal in some cases, under the additional condition that for all primes p and all i we have $U(p, i; A) \geq \aleph_0$ whenever $U(p, i; A) \neq \emptyset$.) The proof of Theorem 1.2 is based

on our definable-type technique from [1] and a detailed analysis of the ordered set of join-irreducibles of the lattice of positive-primitive definable subgroups of A .

No particularly good bounds on density can be expected for 1-based stable theories in general, even in the case of abelian groups (see Lemma 5.22). However, in [1] we already showed that in Presburger Arithmetic, every partitioned formula with n parameter variables has VC density at most n . Theorem 1.2 is the cornerstone of a characterization of those abelian groups admitting a uniform bound on VC density of formulas in terms of the number of parameter variables:

Theorem 1.3. *Let A be an infinite abelian group. Then the following are equivalent:*

1. *there is some d such that each \mathcal{L} -formula $\varphi(x; y)$ with $|y| = 1$ has VC density at most d in A ;*
2. *there is some d such that each \mathcal{L} -formula $\varphi(x; y)$ has VC density at most $d|y|$ in A ;*
3. *there are only finitely many p such that $A[p]$ or A/pA is infinite, and for all p the set $U_{\geq \aleph_0}(p; A)$ is finite.*

Implicit in the arguments comprising the proof of this theorem is an explicit description of all dp-minimal abelian groups. For example, by Theorem 1.2, every \mathcal{L} -formula $\varphi(x; y)$ with $|y| = 1$ has VC density at most 1 in

$$A = \mathbb{Z}(p)^{(\aleph_0)} \oplus \mathbb{Z}(p^2)^{(\aleph_0)},$$

hence this group A is dp-minimal by [9, Proposition 3.2] (see also Section 2.5 below). We refer to Proposition 5.27 for the complete list of all dp-minimal abelian groups.

1.5 Organization of the paper. We begin with a preliminary Section 2, where we recall some basic definitions and results from Part I of our paper. The reader familiar with [1] may choose to skip this section upon first reading. In Section 3 we prove Theorem 1.1 (in a slightly more general form, cf. Theorem 3.1). In Section 4 we introduce the breadth of a module (defined via its lattice of positive-primitive definable subgroups) and relate it to VC density, and after these foundations have been laid, we prove Theorems 1.2 and 1.3 in the final Section 5.

1.6 Notations and conventions. Throughout this paper, d, k, m and n range over the set $\mathbb{N} := \{0, 1, 2, \dots\}$ of natural numbers. We set $[n] := \{1, \dots, n\}$. Given a set X , we write 2^X for the power set of X .

2 Preliminaries

In this section we recapitulate some basic concepts and results proved in the predecessor [1] of this paper. We let \mathbf{M} be an infinite \mathcal{L} -structure for some first-order language \mathcal{L} and $T = \text{Th}(\mathbf{M})$.

2.1 Dual VC density of finite sets of formulas. Let $\Delta = \Delta(x; y)$ be a finite set of partitioned \mathcal{L} -formulas $\varphi = \varphi(x; y)$ in the tuple of object variables x and tuple of parameter variables y . We write $\neg\Delta := \{\neg\varphi : \varphi \in \Delta\}$. Let $B \subseteq M^{|y|}$ be finite. We set

$$\Delta(x; B) := \{\varphi(x; b) : \varphi \in \Delta, b \in B\},$$

and we call a consistent subset of $\Delta(x; B) \cup \neg\Delta(x; B)$ a $\Delta(x; B)$ -type. The set of realizations in \mathbf{M} of a $\Delta(x; B)$ -type p is denoted by $p^{\mathbf{M}}$. Given $a \in M^{|x|}$ we write $\text{tp}^{\Delta}(a/B)$ for the $\Delta(x; B)$ -type realized by a , and

$$S^{\Delta}(B) = \{ \text{tp}^{\Delta}(a/B) : a \in M^{|x|} \}.$$

If $\Delta = \{\varphi\}$ is a singleton, we also write $S^{\varphi}(B)$ instead of $S^{\Delta}(B)$. We set

$$\pi_{\Delta}^*(t) := \max \{ |S^{\Delta}(B)| : B \subseteq M^{|y|}, |B| = t \} \quad \text{for each } t \in \mathbb{N}.$$

If we pass from \mathbf{M} to an elementarily equivalent \mathcal{L} -structure, then π_{Δ}^* does not change; this justifies our notation, which suppresses \mathbf{M} . (But π_{Δ}^* does depend on $T = \text{Th}(\mathbf{M})$, which is also suppressed in our notation.) If no $\varphi \in \Delta$ has the independence property, then there exists a real number r with $0 \leq r \leq \sum_{\varphi \in \Delta} \text{vc}^*(\varphi)$ and

$$|S^{\Delta}(B)| = O(|B|^r) \quad \text{for all finite } B \subseteq M^{|y|}.$$

(See [1, Lemma 3.15].) Hence in this case one can define the *dual VC density* of Δ as

$$\text{vc}^*(\Delta) = \inf \{ r \geq 0 : \pi_{\Delta}^*(t) = O(t^r) \}.$$

If $\Delta = \{\varphi\}$ is a singleton, then $\text{vc}^*(\Delta)$ agrees with $\text{vc}^*(\varphi)$ as defined in the introduction.

2.2 The VC density function of T . Suppose T is NIP. We define the *VC density* of T to be the function $\text{vc} = \text{vc}^T : \mathbb{N} \rightarrow \mathbb{R}^{\geq 0} \cup \{\infty\}$ given by

$$\begin{aligned} \text{vc}(n) &:= \sup \{ \text{vc}(\varphi) : \varphi(x; y) \text{ is an } \mathcal{L}\text{-formula with } |y| = n \} \\ &= \sup \{ \text{vc}^*(\varphi) : \varphi(x; y) \text{ is an } \mathcal{L}\text{-formula with } |x| = n \}. \end{aligned}$$

It is easy to see that $\text{vc}(m) \geq m$ for each m ; in [1, Section 3] we showed that in fact $\text{vc}(m+1) \geq \text{vc}(m) + 1$ for each m . Moreover, we proved (cf. [1, Lemma 3.6]):

Lemma 2.1. *Let $\mathcal{L}' = \mathcal{L} \cup \{c_i : i \in I\}$ where the c_i are new constant symbols, and let $T' \supseteq T$ be a complete \mathcal{L}' -theory. Then $\text{vc}^T = \text{vc}^{T'}$.*

The VC density function of T also bounds the dual VC density of finite sets of formulas; in fact, we have (cf. [1, Corollary 3.19]):

Lemma 2.2. *Let Φ be a set of \mathcal{L} -formulas with the tuple of object variables x and varying parameter variables such that every \mathcal{L} -formula $\varphi(x; y)$ is equivalent in T to a Boolean combination of formulas in Φ . Then*

$$\text{vc}^T(m) = \sup \{ \text{vc}^*(\Delta) : \Delta \subseteq \Phi \text{ finite} \} \quad \text{where } m = |x|.$$

2.3 Breadth. Let \mathcal{B} be a set system on a set X . The smallest integer d (if it exists) with the property that for every non-empty intersection $\bigcap_{i \in [n]} B_i$ of $n > d$ sets from \mathcal{B} there is some $I \subseteq [n]$ with $|I| = d$ and $\bigcap_{i \in [n]} B_i = \bigcap_{i \in I} B_i$, is called the *breadth* of \mathcal{B} ; if there is no such d we also say that \mathcal{B} has breadth ∞ . See Section 2.4 of [1] for more on this concept; the following fact was also already noted there:

Lemma 2.3. *Let G be a group and let \mathcal{H} be a collection of subgroups of G with breadth d . Let $\mathcal{B} = \{gH : g \in G, H \in \mathcal{H}\}$ be the set of all (left) cosets of subgroups from \mathcal{H} . Then \mathcal{B} also has breadth d .*

Proof Let $H_1, \dots, H_n \in \mathcal{H}$ and $g_1, \dots, g_n \in G$ such that $\bigcap_{i \in [n]} g_i H_i \neq \emptyset$. This intersection is a coset of $H := \bigcap_{i \in [n]} H_i$. Choose $I \subseteq [n]$ of size d such that $H = \bigcap_{i \in I} H_i$. Then both $\bigcap_{i \in [n]} g_i H_i$ and $\bigcap_{i \in I} g_i H_i$ are cosets of H , with the former contained in the latter, and hence they are equal. \square

2.4 The VC d property. In the following we let $\Delta = \Delta(x; y)$ be a finite set of partitioned \mathcal{L} -formulas in the object variables x and parameter variables y . The VC d property is a convenient condition on a theory which allows the counting of the number of $\Delta(x; B)$ -types over finite parameter sets B . The definition of the VC d property rests on a “uniform” variant of the notion of definable type, originating in [13]:

Definition 2.4. We say that Δ has *uniform definability of types over finite sets* (abbreviated as *UDTFS*) in \mathbf{M} with d parameters if there are finitely many families

$$\mathcal{D}_i = (\mathbf{d}_{\varphi, i}(y; y_1, \dots, y_d))_{\varphi \in \Delta} \quad (i \in I)$$

of \mathcal{L} -formulas (with $|y_j| = |y|$ for $j = 1, \dots, d$) such that for every finite set $B \subseteq M^{|y|}$ and $q \in S^\Delta(B)$ there are $b_1, \dots, b_d \in B$ and some $i \in I$ such that $\mathcal{D}_i(y; b_1, \dots, b_d)$ defines q . We call the family $\mathcal{D} = (\mathcal{D}_i)_{i \in I}$ a *uniform definition of $\Delta(x; B)$ -types over finite sets* in \mathbf{M} with d parameters.

An easy way to check for UDTFS is via breadth (cf. [1, Lemma 5.2]):

Lemma 2.5. Suppose the set system

$$\mathcal{S}_\Delta = \{\varphi(M^{|x|}; b) : b \in M^{|y|}\}$$

has breadth d . Then Δ has UDTFS with d parameters.

We say that \mathbf{M} has the VC d property if any Δ with $|x| = 1$ has uniform definability of types over finite sets with d parameters. Clearly if \mathbf{M} has the VC d property, then so does every other model of T , and we say that T has the VC d property if some model of T does. The point of the VC d property is that it guarantees that every finite subset of partitioned \mathcal{L} -formulas has UDTFS in \mathbf{M} with a uniform bound on the number of parameters:

Theorem 2.6 (see [1, Theorem 5.7]). Suppose that \mathbf{M} has the VC d property. Then every Δ has UDTFS in \mathbf{M} with $d|x|$ parameters.

If Δ has UDTFS in \mathbf{M} with d parameters, then clearly $|S^\Delta(B)| = O(|B|^d)$ for every finite $B \subseteq M^{|y|}$. Hence if \mathbf{M} has the VC d property, then for each Δ we have $\pi_\Delta^*(t) = O(t^{d|x|})$ for each t and thus $\text{vc}^*(\Delta) \leq d|x|$.

The following fact (cf. [1, Corollary 5.6]), when combined with a quantifier-simplification result, is handy for verifying the VC d property; it will later be applied in the case of modules.

Lemma 2.7. Let Φ be a family of partitioned \mathcal{L} -formulas in the single object variable x such that

1. every partitioned \mathcal{L} -formula in the object variable x is equivalent in T to a Boolean combination of formulas from Φ , and
2. every finite set of \mathcal{L} -formulas from Φ has UDTFS in T with d parameters.

Then T has the VC d property.

2.5 Dp-minimality. By a *bigraph* we mean a triple $G = (X, Y, \Phi)$ where X, Y are sets and $\Phi \subseteq X \times Y$. We call the elements of the disjoint union $X \cup Y$ the *vertices* of the bigraph G and the elements of Φ the *edges* of G . A bigraph $G' = (X', Y', \Phi')$ is a *sub-bigraph* of the bigraph $G = (X, Y, \Phi)$ if $X' \subseteq X$, $Y' \subseteq Y$, and $\Phi' \subseteq \Phi$. We say that a bigraph G contains a given bigraph G' (as a sub-bigraph) if G' is isomorphic to a sub-bigraph of G .

Let $\varphi(x; y)$ be a partitioned \mathcal{L} -formula. Given subsets A, A' of $M^{|x|}$ we define a bigraph $G_{A, A', \varphi} = (A, A', \Phi)$ with edge set

$$\Phi = \{(a, a') \in A \times A' : \{a, a'\} \in \mathcal{S}_\varphi \cap (A \cup A')\}.$$

Note that a pair $(a, a') \in A \times A'$ is connected by an edge of $G_{A, A', \varphi}$ iff

$$\{\varphi^*(y; a), \varphi^*(y; a')\} \cup \{\neg\varphi^*(y; a_0) : a_0 \in A \cup A', a_0 \neq a, a'\}$$

is realized in \mathbf{M} .

It is easy to see that a bigraph $G = (X, Y, \Phi)$ with a finite number n of vertices can have at most $\frac{1}{4}n^2$ edges. A fundamental fact about bigraphs is the theorem of Kővári, Sós and Turán [22]: *given a positive integer r , there exists a real number $C = C(r)$ such that every bigraph G with n vertices which does not contain $K_{r,r}$ as a sub-bigraph has at most $Cn^{2-1/r}$ edges.* Here $K_{r,r}$ denotes the (complete) bigraph $K_{r,r} = ([r], [r], [r] \times [r])$. In particular, from this theorem we obtain:

Lemma 2.8. *Let $\varphi(x; y)$ be a partitioned \mathcal{L} -formula. The following are equivalent:*

1. *for each $r > 0$ there are finite subsets A, A' of $M^{|x|}$ such that $G_{A, A', \varphi}$ contains the bigraph $K_{r,r}$;*
2. *for each $r > 0$ there are finite subsets A, A' of $M^{|x|}$ such that $G_{A, A', \varphi}$ has more than $Cn^{2-1/r}$ edges, where $C = C(r)$ is as above and $n = |A| + |A'|$.*

The \mathcal{L} -structure \mathbf{M} is said to be *dp-minimal* if there is no \mathcal{L} -formula $\varphi(x; y)$ with $|y| = 1$ satisfying one of the two equivalent conditions in the previous lemma. In [1] we called \mathbf{M} *vc-minimal* if $\text{vc}^*(\varphi) < 2$ for every \mathcal{L} -formula $\varphi(x; y)$ with $|x| = 1$. (So if $\text{vc}(1) < 2$, then \mathbf{M} is vc-minimal.) It is easy to see (cf. [1, Corollary 5.13]) that if \mathbf{M} is vc-minimal, then \mathbf{M} is dp-minimal. (We do not know whether, conversely, every dp-minimal structure is vc-minimal.) On the other hand, if \mathbf{M} is dp-minimal, then for each \mathcal{L} -formula $\varphi(x; y)$ with $|y| = 1$ there is an upper bound on the size of finite subsets $A \subseteq M^{|x|}$ such that $A \cap \mathcal{S}_\varphi$ contains all two-element subsets of A . Thus, although reducts of dp-minimal structures are obviously also dp-minimal, a structure interpretable in a dp-minimal structure is not in general itself dp-minimal:

Example. Suppose \mathcal{L} is the empty language, and let $\mathcal{L}' = \{\pi_1, \pi_2\}$ where π_i is a unary function symbol, for $i = 1, 2$. Let \mathbf{M}' be the \mathcal{L}' -structure with underlying set $M' = M \times M$, where $\pi_i^{M'}$ is given by $(x_1, x_2) \mapsto (x_i, x_i)$, for $i = 1, 2$. Then clearly \mathbf{M}' is interpretable without parameters in M . Consider the partitioned \mathcal{L}' -formula $\varphi(x; y) := x = \pi_1(y) \vee x = \pi_2(y)$, where $|x| = |y| = 1$. Then $\mathcal{S}_\varphi^{M'} = \binom{D}{1} \cup \binom{D}{2}$, where $D = \{(a, a) : a \in M\}$, hence \mathbf{M}' is not dp-minimal.

One defines a complete \mathcal{L} -theory (without finite models) to be dp-minimal if one of its models is dp-minimal (equivalently, if all of its models are). The definition of dp-minimality given here is not the original one as introduced in [28], which used the notion of ICT pattern from [39]: an *ICT pattern* in \mathbf{M} consists of a pair

$\alpha(x; y)$, $\beta(x; y)$ of partitioned \mathcal{L} -formulas, where $|x| = 1$, and sequences $(a_i)_{i \in \mathbb{N}}$, $(b_j)_{j \in \mathbb{N}}$ in $M^{|y|}$ such that for all i and j the set

$$\{\alpha(x; a_i), \beta(x; b_j)\} \cup \{\neg\alpha(x; a_k) : k \neq i\} \cup \{\neg\beta(x; b_l) : l \neq j\}$$

is consistent with \mathbf{M} . One can show (see, e.g., [9, Fact 2.13] or [1, Proposition 5.12]) that \mathbf{M} is dp-minimal iff there is no ICT pattern in an elementary extension of \mathbf{M} .

The extent of the failure of the dp-minimality property for a theory can be quantified via its dp-rank.

Definition 2.9. One says that T has dp-rank at least n if there are partitioned \mathcal{L} -formulas $\varphi_1(x; y), \dots, \varphi_n(x; y)$, where $|x| = 1$, and parameters $a_j^i \in M^{|y|}$, where $i \in [n]$ and $j \in \mathbb{N}$, so that for any function $\eta: [n] \rightarrow \mathbb{N}$ the set of $\mathcal{L}(M)$ -formulas

$$\{\varphi_i(x; a_{\eta(i)}^i) : i \in [n]\} \cup \{\neg\varphi_i(x; a_j^i) : i \in [n], j \neq \eta(i)\}$$

is consistent (with \mathbf{M}).

So T is not dp-minimal iff T has dp-rank at least 2. An easy generalization of [1, Corollary 5.13] shows that if T has dp-rank at least n , then $\text{vc}(1) \geq n$. The following lemma gives a convenient way to check for dp-rank at least n :

Lemma 2.10. *Suppose there exists an infinite definable set $X \subseteq M$ and a definable one-to-one function $X^n \rightarrow M$. Then the dp-rank of $T = \text{Th}(\mathbf{M})$ is at least n .*

Proof Let $f: X^n \rightarrow M$ be definable and injective, where $X \subseteq M$ is definable. Without loss of generality we may assume that X is \emptyset -definable by $\psi(x)$ and f is also \emptyset -definable. Pick pairwise distinct $a_i \in X$ for $i \in \mathbb{N}$. Let $\varphi_i(x; y)$ be the formula

$$\exists z_1 \dots z_n \left(\bigwedge_{j \in [n]} \psi(z_j) \wedge f(z_1, \dots, z_n) = x \wedge z_i = y \right).$$

Now as in the preceding definition setting $a_j^i := a_j$ for $i \in [n]$ and $j \in \mathbb{N}$, the formulas φ_i and parameters a_j^i witness that T has dp-rank at least n . \square

3 Linear VC Density for Finite Rank Theories

Throughout this section we assume that T is a complete stable theory without finite models, in a first-order language \mathcal{L} . We fix a monster model \mathbf{M} of T . “Small” means “of cardinality smaller than $|M|$,” and “ $\models \dots$ ” abbreviates “ $\mathbf{M} \models \dots$ ” We also let Δ be a set (possibly infinite) of partitioned \mathcal{L} -formulas in the tuple of object variables $x = (x_1, \dots, x_m)$ (and varying tuples of parameter variables).

Given $B \subseteq M$, by a Δ -formula over B we mean a formula of the form $\varphi(x; b)$ or $\neg\varphi(x; b)$ where $\varphi(x; y) \in \Delta$ and $b \in B^{|y|}$, and by a Δ -type over B (in \mathbf{M}) we mean a consistent set of Δ -formulas over B . A Δ -type $p = p(x)$ over B is said to be complete if for all Δ -formulas φ , either φ or $\neg\varphi$ is in p . We denote by $S_\Delta(B)$ the set of complete Δ -types over B in \mathbf{M} , equipped with the topology with a subbasis consisting of the sets of the form $\{p \in S_\Delta(B) : \varphi \in p\}$ where φ is a Δ -formula over B . This makes $S_\Delta(B)$ into a compact totally disconnected Hausdorff space. As usual $S_x(B)$ (or $S_m(B)$, if the particular choice of an m -tuple x of object variables is unimportant) denotes $S_\Delta(B)$ where Δ is the set of all partitioned \mathcal{L} -formulas in the object variables x . (Note the difference between the notions $S_\Delta(B)$ and

$S^\Delta(B)$ as introduced in Section 2 for finite sets of \mathcal{L} -formulas $\Delta = \Delta(x; y)$ in the parameter variables y : in the former the parameter set B is understood to be a subset of M , in the latter it is a subset of $M^{|y|}$.)

We will write $R_\Delta(-)$ for the rank denoted by $R^m(-, \Delta, \aleph_0)$ in [38, Definition II.1.1]. Given a set $\Theta = \Theta(x)$ of $\mathcal{L}(M)$ -formulas, denote by $[\Theta]$ the closed subset of $S_\Delta(M)$ consisting of all $p(x) \in S_\Delta(M)$ which are consistent with $\Theta(x)$. Then $R_\Delta(\Theta)$ is the Cantor-Bendixson rank of $[\Theta]$ (i.e., the maximum Cantor-Bendixson rank of an element of $[\Theta]$), with the convention that $R_\Delta(\Theta) = -1$ if Θ is inconsistent, and if $[\Theta]$ does not have a Cantor-Bendixson rank, then $R_\Delta(\Theta) := \infty$. We call $R_\Delta(\Theta)$ the Δ -rank of Θ . If $\Theta = \{\theta\}$ is a singleton we also write $R_\Delta(\theta)$ instead of $R_\Delta(\Theta)$. We also let $R_\Delta(T) := R_\Delta(x = x)$ where $x = x$ is shorthand for the \mathcal{L} -formula $x_1 = x_1 \wedge \dots \wedge x_m = x_m$. Since we assume that T is stable, for finite Δ , the rank $R_\Delta(-)$ always takes values smaller than ω [38, Lemma II.2.1 and Theorem II.2.2].

Our goal in this section is to prove:

Theorem 3.1. *Suppose T does not have the finite cover property and Δ is finite, and let $\varrho = R_\Delta(T)$. Then $\pi_\Delta^*(t) = O(t^\varrho)$ and hence $\text{vc}^*(\Delta) \leq \varrho$.*

See Corollary 3.8 below on how this theorem implies Theorem 1.1 from the introduction. Before we give the proof of Theorem 3.1 we recall a few facts about R_Δ used in the proof. We will occasionally work in M^{eq} . Nonetheless, all sets or elements should be assumed to lie in the home sort unless explicitly stated otherwise.

3.1 Basic properties of Δ -rank. Let, as above, Δ be a set of partitioned \mathcal{L} -formulas in the object variables $x = (x_1, \dots, x_m)$, and let $\Theta(x)$ be a set of $\mathcal{L}(M)$ -formulas.

If $R_\Delta(\Theta) < \infty$, then there is a largest integer d such that there are pairwise distinct $p_1, \dots, p_d \in [\Theta]$ of maximal Cantor-Bendixson rank $R_\Delta(\Theta)$; we call d the Δ -degree of Θ , denoted by $\deg_\Delta(\Theta)$. (In [38, Definition II.1.1] this is denoted by $\text{Mlt}(\Theta, \Delta, \aleph_0)$.) For an $\mathcal{L}(M)$ -formula θ we write $\deg_\Delta(\theta) := \deg_\Delta(\{\theta\})$. For every type $\Theta(x)$, there is a $p \in S_x(M)$ with $p \supseteq \Theta$ and $R_\Delta(p) = R_\Delta(\Theta)$; there are exactly $\deg_\Delta(\Theta)$ many types in $S_\Delta(M)$ of the form $p \upharpoonright \Delta$ where $p \in S_x(M)$ with $p \supseteq \Theta$ and $R_\Delta(p) = R_\Delta(\Theta)$. If Θ is closed under finite conjunctions, then there is a formula $\theta \in \Theta$ such that $R_\Delta(\Theta) = R_\Delta(\theta)$ and $\deg_\Delta(\Theta) = \deg_\Delta(\theta)$.

If $a \in M^m$ and $B \subseteq M^n$, then we set $R_\Delta(a/B) := R_\Delta(\text{tp}(a/B'))$ where B' is the smallest subset of M such that $B \subseteq (B')^n$, and similarly for \deg_Δ . (Note that $|B'| \leq n|B|$.) If $\Theta(x)$ is a type over a small set $B \subseteq M$, then there exists a realization $a \in M^{|x|}$ of Θ with $R_\Delta(a/B) = R_\Delta(\Theta)$; we say that such a realization of Θ is Δ -generic.

If Δ' is another set of partitioned \mathcal{L} -formulas in the object variables x with $\Delta \subseteq \Delta'$, then $R_\Delta(\Theta) \leq R_{\Delta'}(\Theta)$, and if equality holds, then $\deg_\Delta(\Theta) \leq \deg_{\Delta'}(\Theta)$ [38, Lemma II.1.3]. If Δ is the set of all partitioned \mathcal{L} -formulas in the tuple of object variables x , then R_Δ is Morley rank; hence $R_\Delta(\Theta) \leq \text{MR}(\Theta)$. Moreover, if Θ' is another set of $\mathcal{L}(M)$ -formulas with $\Theta^M \subseteq (\Theta')^M$, then $R_\Delta(\Theta) \leq R_\Delta(\Theta')$, and if equality holds, then $\deg_\Delta(\Theta) \leq \deg_{\Delta'}(\Theta')$ [38, Theorem II.1.1]. It is also easy to see that for all $\mathcal{L}(M)$ -formulas $\theta_1(x)$ and $\theta_2(x)$ we have

$$R_\Delta(\theta_1 \vee \theta_2) = \max\{R_\Delta(\theta_1), R_\Delta(\theta_2)\}. \quad (3.1)$$

Suppose now $A \subseteq B \subseteq M$, $p \in S_x(B)$. Then $R_\Delta(p \upharpoonright A) \geq R_\Delta(p)$ for all Δ , and p does not fork over A iff $R_\Delta(p \upharpoonright A) = R_\Delta(p)$ for all finite Δ [38, Theorem III.4.1]. For example, if $B = \text{acl}(A)$, then p does not fork over A (since any two extensions of a type in $S_x(A)$ to types in $S_x(\text{acl}(A))$ are conjugate under $\text{Aut}(M \upharpoonright A)$ and hence have the same Δ -rank). A complete type is stationary (i.e., has a unique non-forking extension to any larger parameter set) iff it has Δ -degree 1 for all finite Δ [38, Theorem III.4.2], and as a consequence of the Finite Equivalence Relation Theorem [38, Theorem III.2.8], every complete type over a parameter set which is algebraically closed in M^{eq} is stationary; hence we obtain:

Lemma 3.2. *Suppose Δ is finite, and let $\theta(x)$ be an $\mathcal{L}(A)$ -formula, $A \subseteq M$, and set $\alpha = R_\Delta(\theta)$, $d = \deg_\Delta(\theta)$. There exist $\mathcal{L}^{\text{eq}}(\text{acl}^{\text{eq}}(A))$ -formulas $\theta_1(x), \dots, \theta_d(x)$ such that $R_\Delta(\theta_i) = \alpha$, $\deg_\Delta(\theta_i) = 1$ for every $i \in [d]$, and $\theta(M^m)$ is the disjoint union of the sets $\theta_i(M^m)$, $i \in [d]$.*

We finish the preliminaries by recording three important consequences of a theory not having the finite cover property.

Fact 3.3. [38, Theorem II.4.4] Suppose T does not have the finite cover property and Δ is finite. Then:

1. There is $k \in \mathbb{N}$ so that if p is any Δ -type there is $q \subseteq p$ with $|q| \leq k$ so that $R_\Delta(p) = R_\Delta(q)$. (Finite witness of rank.)
2. For any \mathcal{L} -formula $\theta(x; y)$ and any $\alpha \in \mathbb{N}$ there is an \mathcal{L} -formula $\psi(y)$ so that for any $b \in M^{|y|}$, $R_\Delta(\theta(x; b)) = \alpha$ iff $\models \psi(b)$. (Definability of rank.)
3. For any \mathcal{L} -formula $\theta(x; y)$ there is some $D \in \mathbb{N}$ so that for any $b \in M^{|y|}$, $\deg_\Delta(\theta(x; b)) \leq D$. (Boundedness of multiplicity.)

3.2 Proof of Theorem 3.1. The proof will be through a series of lemmas. We fix a finite set $\Delta = \Delta(x; y)$ of partitioned \mathcal{L} -formulas, and write $\varrho = R_\Delta(T)$.

Lemma 3.4. *Let $\psi(x) = \bigwedge_{i=1}^k \varphi_i(x; b_i)$ be a conjunction of instances of formulas $\varphi_i(x; y_i) \in \Delta$, $\alpha := R_\Delta(\psi)$, $\mu := \min_i R_\Delta(\varphi_i(x; b_i))$, and $a \in M^m$ with $\models \psi(a)$ and $R_\Delta(a/b_1, \dots, b_k) = \alpha$. There is an \mathcal{L}^{eq} -formula $\theta(x; z)$ and $c \in \text{acl}^{\text{eq}}(b_{i_1}, \dots, b_{i_r})$, where $i_1, \dots, i_r \in [k]$ with $r \in [\mu - \alpha]$, such that $\models \theta(a; c)$ and $R_\Delta(\theta(x; c)) = \alpha$.*

Proof Let $\gamma_1(x; c_1)$, where $c_1 \in \text{acl}^{\text{eq}}(\emptyset)$, have Δ -rank $\mu_1 := \varrho$ and Δ -degree 1 so that $\models \gamma_1(a; c_1)$, by Lemma 3.2. We have $\mu_1 \geq \alpha$; if $\mu_1 = \alpha$ we are done, so suppose $\mu_1 > \alpha$. Then $R_\Delta(\gamma_1(x; c_1) \wedge \varphi_i(x; b_i)) < \mu_1$ for some i : otherwise $R_\Delta(\gamma_1(x; c_1) \wedge \neg \varphi_i(x; b_i)) < \mu_1$ for every i , since $\deg_\Delta(\gamma_1(x; c_1)) = 1$, hence $R_\Delta(\gamma_1(x; c_1) \wedge \neg(\bigwedge_i \varphi_i(x; b_i))) < \mu_1$ and therefore $R_\Delta(\gamma_1(x; c_1)) < \mu_1$, by (3.1), a contradiction. Without loss of generality we may assume that $\mu_2 := R_\Delta(\gamma_1(x; c_1) \wedge \varphi_1(x; b_1)) < \mu_1$. We have $\mu_2 \geq \alpha$, and if $\mu_2 = \alpha$ we are done. Otherwise let $\gamma_2(x; c_2)$, where $c_2 \in \text{acl}^{\text{eq}}(b_1)$, have Δ -rank μ_2 and Δ -degree 1 so that $\models \gamma_2(a; c_2)$ and $\models \forall x(\gamma_2(x; c_2) \rightarrow \gamma_1(x; c_1) \wedge \varphi_1(x; b_1))$, according to Lemma 3.2. As before we see that for some $i > 1$ we have $R_\Delta(\gamma_2(x; c_2) \wedge \varphi_i(x; b_i)) < \mu_2$.

We can continue in this fashion to find a sequence $\gamma_1(x; z), \dots, \gamma_r(x; z)$ of \mathcal{L}^{eq} -formulas, a sequence of indices $i_1, \dots, i_r \in [k]$, and a sequence c_1, \dots, c_r where $c_j \in \text{acl}^{\text{eq}}(b_{i_1}, \dots, b_{i_{j-1}})$, of maximal length $r \geq 1$, subject to the following properties: for every $j \in [r]$ we have $\models \gamma_j(a; c_j)$, and setting $\delta_0(x) := (x = x)$ and

$\delta_j(x; y_{i_j}, z) := \gamma_j(x; z) \wedge \varphi_{i_j}(x; y_{i_j})$ for $j \in [r]$, we have:

$$\models \forall x(\gamma_j(x; c_j) \rightarrow \delta_{j-1}(x; b_{i_{j-1}}, c_{j-1}))$$

and

$$\alpha \leq R_\Delta(\delta_j(x; b_{i_j}, c_j)) < R_\Delta(\delta_{j-1}(x; b_{i_{j-1}}, c_{j-1})).$$

We have $r \leq \mu - \alpha$, and by the above $R_\Delta(\delta_r(x; b_{i_r}, c_r)) = \alpha$. Hence

$$\theta(x; y_{i_r}, z) := \delta_r(x; y_{i_r}, z), \quad c := (b_{i_r}, c_r)$$

have the required properties. \square

Lemma 3.5. *Suppose T does not have the finite cover property. Let*

$$\psi(x; y_1, \dots, y_k) = \bigwedge_{i=1}^k \varphi_i(x; y_i) \quad \text{where } \varphi_i(x; y_i) \in \Delta \text{ for } i \in [k],$$

and let $\alpha \in [\varrho]$. There is a family

$$\{\theta_{ij}(x; y_{i_1}, \dots, y_{i_{\varrho-\alpha}}) : \mathbf{i} = (i_1, \dots, i_{\varrho-\alpha}) \in [k]^{\varrho-\alpha}, j \in [N]\}$$

of \mathcal{L} -formulas, where $N \in \mathbb{N}$, with the following properties: if $a \in M^m$ and $b_i \in M^{|y_i|}$ ($i \in [k]$) are such that

$$\models \psi(a; b_1, \dots, b_k) \quad \text{and} \quad \alpha = R_\Delta(a/b_1, \dots, b_k) = R_\Delta(\psi(x; b_1, \dots, b_k)),$$

then for some $\mathbf{i} \in [k]^{\varrho-\alpha}$ and $j \in [N]$ we have

$$\models \theta_{ij}(a; b_{i_1}, \dots, b_{i_{\varrho-\alpha}}) \quad \text{and} \quad R_\Delta(\theta_{ij}(x; b_{i_1}, \dots, b_{i_{\varrho-\alpha}})) = \alpha.$$

Proof For each $a \in M^m$ and $b_i \in M^{|y_i|}$ ($i \in [k]$) satisfying the hypothesis, by Lemma 3.4 there are $i_1, \dots, i_{\varrho-\alpha} \in [k]$ such that

$$R_\Delta(a/\text{acl}^{\text{eq}}(b_{i_1}, \dots, b_{i_{\varrho-\alpha}})) = \alpha$$

and hence that

$$R_\Delta(a/b_{i_1}, \dots, b_{i_{\varrho-\alpha}}) = \alpha,$$

so for some \mathcal{L} -formula $\theta(x; y_{i_1}, \dots, y_{i_{\varrho-\alpha}})$ such that $\models \theta(a; b_{i_1}, \dots, b_{i_{\varrho-\alpha}})$ we have

$$R_\Delta(\theta(x; b_{i_1}, \dots, b_{i_{\varrho-\alpha}})) = \alpha.$$

Thus the set of formulas $\Gamma(z; y_1, \dots, y_k)$, where $|z| = m$, consisting of

- $\psi(z; y_1, \dots, y_k)$;
- $R_\Delta(\psi(x; y_1, \dots, y_k)) = \alpha$;
- $R_\Delta(\theta(x; y_1, \dots, y_k)) < \alpha \rightarrow \neg\theta(z; y_1, \dots, y_k)$, where θ ranges over all partitioned \mathcal{L} -formulas in $(x; y_1, \dots, y_k)$; and
- $\neg(\theta(z; y_{i_1}, \dots, y_{i_{\varrho-\alpha}}) \wedge R_\Delta(\theta(x; y_{i_1}, \dots, y_{i_{\varrho-\alpha}})) = \alpha)$, where $i_1, \dots, i_{\varrho-\alpha}$ range over $[k]$ and θ over all partitioned \mathcal{L} -formulas in $(z; y_{i_1}, \dots, y_{i_{\varrho-\alpha}})$,

is inconsistent. Note that this is a first-order type by Fact 3.3, (2). The lemma now follows by compactness. \square

Lemma 3.6. *Suppose T does not have the finite cover property. There is a family $\{\theta_j(x; y_1, \dots, y_\varrho)\}_{j \in [N]}$ of \mathcal{L} -formulas, where $N \in \mathbb{N}$ and $|y_i| = |y|$ for $i \in [\varrho]$, with the following property: if $B \subseteq M^{|y|}$, $p \in S^\Delta(B)$, and $a \in M^m$ is a Δ -generic realization of p , then there are $b_1, \dots, b_\varrho \in B$ and $j \in [N]$ such that $\models \theta_j(a; b_1, \dots, b_\varrho)$ and*

$$R_\Delta(\theta_j(x; b_1, \dots, b_\varrho)) = R_\Delta(p).$$

Proof We may and shall assume that Δ is closed under negation, i.e., with every $\varphi \in \Delta$ the set Δ also contains a formula equivalent to $\neg\varphi$ (in T). By Fact 3.3, (1) take $k \in \mathbb{N}$ so that if p is any Δ -type, then $R_\Delta(p) = R_\Delta(q)$ for some $q \subseteq p$ with $|q| \leq k$.

Apply Lemma 3.5 to all possible conjunctions of k formulas from Δ (and then combine the resulting families $\{\theta_{ij}\}$ to a single family) to obtain a family $\{\theta_j(x; y_1, \dots, y_\varrho)\}_{j \in [N]}$ of \mathcal{L} -formulas, where $N \in \mathbb{N}$ and $|y_i| = |y|$ for $i \in [\varrho]$, with the following property: if $\psi(x) = \bigwedge_{i=1}^k \varphi_i(x; b_i)$ is a conjunction of instances of formulas from Δ and $a \in M^m$ with $\models \psi(a)$ and $R_\Delta(a/b_1, \dots, b_k) = R_\Delta(\psi(x))$, then for some $j \in [N]$ we have

$$\models \theta_j(a; b_{i_1}, \dots, b_{i_\varrho}) \quad \text{and} \quad R_\Delta(\theta_j(x; b_{i_1}, \dots, b_{i_\varrho})) = R_\Delta(a/b_1, \dots, b_k).$$

Now let $B \subseteq M^{|y|}$, $p \in S^\Delta(B)$, and $a \in M^m$ be a realization of p with $R_\Delta(a/B) = R_\Delta(p)$. Choose $\varphi_i(x; y) \in \Delta$ and $b_i \in B$, where $i \in [k]$, such that $R_\Delta(\bigwedge_{i=1}^k \varphi_i(x; b_i)) = R_\Delta(p)$. Then there are $i \in [k]^\varrho$ and $j \in [N]$ such that $\models \theta_j(a; b_{i_1}, \dots, b_{i_\varrho})$ and $R_\Delta(\theta_j(x; b_{i_1}, \dots, b_{i_\varrho})) = R_\Delta(p)$. \square

Proof of Theorem 3.1. We wish to count the number of $\Delta(x; B)$ -types over finite parameter sets $B \subseteq M^{|y|}$. Fix \mathcal{L} -formulas $\theta_i(x; y_1, \dots, y_\varrho)$ with $i \in [N]$ as in Lemma 3.6. Let $D \in \mathbb{N}$ bound the Δ -degree of any instance of any θ_i . Note that D exists by Fact 3.3, (3). Fix a finite set $B \subseteq M^{|y|}$. Let $p \in S^\Delta(B)$ and let $a \in M^m$ be a Δ -generic realization of p . There are $b_1, \dots, b_\varrho \in B$ and $i \in [N]$ so that $\models \theta_i(a; b_1, \dots, b_\varrho)$ and $R_\Delta(p) = R_\Delta(\theta_i(x; b_1, \dots, b_\varrho))$. Since $\deg_\Delta(\theta_i(x; b_1, \dots, b_\varrho)) \leq D$ there are global Δ -types $q_1, \dots, q_D \in S_\Delta(M)$ (not necessarily distinct) such that for any a' with $\models \theta_i(a'; b_1, \dots, b_\varrho)$ and $R_\Delta(a'/B) = R_\Delta(\theta_i(x; b_1, \dots, b_\varrho))$ we have $\text{tp}^\Delta(a'/B) \subseteq q_i$ some $i \in [D]$. In particular $p \subseteq q_i$ for some $i \in [D]$. So $|S^\Delta(B)| \leq ND|B|^\varrho$, and our result follows. \square

3.3 Δ -rank, U-rank, and Morley rank. We now want to apply Theorem 3.1 to obtain a uniform bound on the VC densities $\text{vc}^*(\Delta)$ of finite sets of \mathcal{L} -formulas $\Delta = \Delta(x; y)$ in terms of $m = |x|$ and the U-rank $U(T)$ of T , provided the latter is finite (i.e., less than ω). For this, we first note the following relationship between the ranks R_Δ and the Lascar U-rank of a complete type:

Lemma 3.7. *Let $p \in S_x(A)$ where $A \subseteq M$. Then:*

1. $R_\Delta(p) \leq U(p)$ for all finite sets $\Delta = \Delta(x; y)$ of partitioned \mathcal{L} -formulas;
2. if $U(p) < \omega$, then there is some finite set $\Delta = \Delta(x; y)$ of partitioned \mathcal{L} -formulas such that $R_\Delta(p) = U(p)$.

Proof For (1) recall that the Δ -ranks detect forking: $q \in S_x(B)$ with $A \subseteq B \subseteq M$ does not fork over A iff $R_\Delta(q \upharpoonright A) = R_\Delta(q)$ for every finite set $\Delta = \Delta(x; y)$ of partitioned \mathcal{L} -formulas. Given such Δ , a straightforward transfinite induction now proves the implication $R_\Delta(p) \geq \alpha \Rightarrow U(p) \geq \alpha$, for all ordinals α , as required.

For (2), we use that by [38, Lemma III.1.2, (2)], if a type $q \in S_x(B)$ with $A \subseteq B \subseteq M$ forks over A , then there is a finite Δ_0 such that for every finite Δ containing Δ_0 one has $R_\Delta(q) < R_\Delta(q \upharpoonright A)$. So from a forking sequence $p = p_0 \subseteq p_1 \subseteq \dots \subseteq p_n$ of types of length n one obtains a finite Δ such that $R_\Delta(p_0) > R_\Delta(p_1) > \dots > R_\Delta(p_n) \geq 0$ and so $R_\Delta(p) = R_\Delta(p_0) \geq n$. \square

A special case of Lascar's Inequality (see, e.g., [34, Theorem 1.5]) states that if $U(T) < \omega$ and $a_1 \in M^{m_1}$ and $a_2 \in M^{m_2}$ are independent (over \emptyset), then $U(a_1, a_2) = U(a_1) + U(a_2)$. (Here $U(a) := U(\text{tp}(a/\emptyset))$ for $a \in M^m$.) Thus in this case all types in $S_m(T) = S_m(\emptyset)$ have U -rank at most $m U(T)$. Hence from Theorem 3.1 and the previous lemma we obtain the following linear bound on the VC density function of T (that is, Theorem 1.1 from the introduction):

Corollary 3.8. *Suppose T has finite U -rank $d = U(T)$ and does not have the finite cover property. Then if Δ is finite we have $\pi_\Delta^*(t) = O(t^{d|x|})$, and hence $\text{vc}^T(m) \leq dm$ for each m .*

Since $U(T) \leq \text{MR}(T)$, the previous corollary applies in particular if T has finite Morley rank and does not have the finite cover property. The following folklore result (see, e.g., [29, Proposition B.1], [31], or [34, Theorem 4.5]) can be used to verify that T has finite Morley rank:

Proposition 3.9. *Suppose there is a family $\{D_i\}_{i \in I}$ of strongly minimal definable sets such that every non-algebraic type in T is non-orthogonal to some D_i . Then $U(p) = \text{MR}(p) < \omega$ for each $p \in S_m(B)$, $B \subseteq M$.*

Example 3.10. Suppose T is totally transcendental and unidimensional (i.e., any two non-algebraic 1-types are non-orthogonal). Then T satisfies the hypothesis of the previous proposition with a single D_i . (The fact that $\text{MR}(T)$ is finite for \aleph_1 -categorical countable T was first shown in [3].)

3.4 Applications. Together with the remarks in the previous subsection, we can now draw some immediate consequences of Theorem 3.1.

3.4.1 *Totally transcendental \aleph_0 -categorical theories.* First we discuss the \aleph_0 -categorical case. It is well-known that a stable \aleph_0 -categorical theory does not have the finite cover property. (This is a consequence of [38, Theorem II.4.4], see [34, Lemma 2.4].) If T is countable, \aleph_0 -categorical, and ω -stable, then T has finite Morley rank, in fact, $\text{MR}(T) \leq |S_2(\emptyset)|$; see [7, Corollary 4.3 and Theorem 5.1]. In this case moreover $U(p) = \text{MR}(p)$ for each m -type p over a finite parameter set, by [34, Theorem 2.1]. Hence by Corollary 3.8:

Corollary 3.11. *Suppose T is countable, totally transcendental, and \aleph_0 -categorical. Then $\text{vc}^T(m) \leq m U(T) = m \text{MR}(T) \leq m |S_2(\emptyset)|$ for every m .*

3.4.2 *Dimensional theories.* Recall that T is called dimensional (also known as non-multidimensional) if the set of domination equivalence classes of non-algebraic stationary 1-types is small. Dimensional theories do not have the finite cover property (see [38, Lemma IX.1.10]). Hence by Corollary 3.8 and Example 3.10 above:

Corollary 3.12. *Suppose T is totally transcendental and unidimensional. Then $U(T) = \text{MR}(T) < \omega$ and $\text{vc}^T(m) \leq m \text{MR}(T)$ for every m .*

The previous corollary applies if T is countable and \aleph_1 -categorical. In particular, we see that if T is strongly minimal, then $\text{vc}^T(m) = m$ for every m . For example, if $T = \text{ACF}_p$ is the theory of algebraically closed fields of characteristic p (where p is a prime or $p = 0$), then $\text{vc}^T(m) = m$ for every m .

Every U -rank 1 theory is dp-minimal [28, Fact 3.2]. Our next application strengthens this remark:

Corollary 3.13. *Suppose $U(T) = 1$. Then $\text{vc}^T(m) = m$ for every m .*

Proof By Theorem 21 in [5], after expanding \mathcal{L} by constant symbols if necessary, one may assume that T is *quasi strongly minimal* in the sense of [5], that is: in each model of T , the one-variable definable sets are precisely the Boolean combinations of \emptyset -definable or finite sets. (Quasi strongly minimal theories are the analogue, in stability theory, of quasi weakly o-minimal theories.) In this case, by [5, Theorem 20] every non-algebraic 1-type $p \in S_1(A)$ is determined by its restriction $p \upharpoonright \emptyset$ to the empty parameter set. In particular, T is dimensional, and hence does not have the finite cover property, thus Corollary 3.8 applies. \square

For a simple example where this corollary applies (which can also be dealt with directly), consider the language $\mathcal{L} = \{P_i : i \in \mathbb{N}\}$, where each P_i is a unary predicate symbol, and the \mathcal{L} -theory T of an infinite set with the P_i interpreted by independent predicates, i.e.:

$$T \models \exists^{\geq n} x \left(\bigwedge_{i \in I} P_i x \wedge \bigwedge_{j \in J} \neg P_j x \right) \quad \text{for all } n \text{ and finite disjoint } I, J \subseteq \mathbb{N}.$$

(See also [9, Proposition 3.7].)

3.4.3 Groups of finite Morley rank. Groups of finite Morley rank are finite-dimensional and hence do not have the finite cover property [23, Lemme 4, p. 457]. Hence by Corollary 3.8:

Corollary 3.14. *Suppose \mathcal{L} is an expansion of the language of groups and T is an expansion of the theory of groups with $\text{MR}(T) < \omega$. Then $\text{vc}^T(m) \leq m \text{MR}(T)$ for every m .*

This corollary applies, in particular, to the complete theory of a module of finite Morley rank. However, in this case we may use instead Corollary 4.15 and Lemma 4.25 below to obtain the same result.

3.5 Examples and remarks. We first point out that the bound on VC density given in Theorem 3.1 is far from sharp:

Example. Let $\mathcal{L} = \{E_n : n \geq 0\}$ be the language consisting of countably many binary relations E_n . Let T be the theory whose axioms state that each E_n is an equivalence relation with infinitely many infinite classes and so that E_{n+1} refines E_n in such a way that each E_n -class is partitioned into infinitely many E_{n+1} -classes. The theory T is stable and does not have the finite cover property. Fix some n and let $\Delta(x; y) = \{E_0(x; y), \dots, E_{n-1}(x; y)\}$. It is easy to see that $\text{vc}^*(\Delta) \leq 1$ yet $R_\Delta(x = x) \geq n$.

Similarly one may ask if not having the finite cover property is necessary for T to have minimal density. This is also seen to be false by a simple example.

Example. Let \mathcal{L} be the language with a single binary relation E and let T be the theory stating that E is an equivalence relation with one class of size n for each n . Then T is ω -stable and has the finite cover property. But it is easy to see that $\text{vc}^*(\Delta) \leq |x|$ for any finite set $\Delta = \Delta(x; y)$ of partitioned \mathcal{L} -formulas.

Laskowski [25] has shown that for every \mathcal{L} -formula $\varphi(x; y)$ there is an \mathcal{L} -formula $d_\varphi(y; y_1, \dots, y_d)$ with $d = R^m(x = x, \varphi, 2)$, where $m = |x|$, so that if B is a set of $|y|$ -tuples (possibly infinite) and $p \in S_\varphi(B)$, then there are $b_1, \dots, b_d \in B$ so




Figure 3.1 An E -equivalence class C in the case $n = 2$

that $d_\varphi(y; b_1, \dots, b_d)$ defines p , and hence $\text{vc}^*(\varphi) \leq d$. (See also [1, Section 5.3].) We may ask if a similar situation holds in the case that T does not have the finite cover property and we replace $R^m(-, \varphi, 2)$ by $R_\varphi(-)$. If this were the case, then our analysis of VC density in the superstable finite-rank case would mirror much of our analysis in other cases given in [1], in that bounds on VC density arise out of a bound on the number of parameters needed to define types. We show by a very simple example that this is not the case. Namely we show that for any $n > 0$ there is a totally categorical theory T_n with $\text{MR}(T_n) = 1$ in which for a fixed formula $\varphi(x; y)$ with $|x| = 1$ the defining scheme for φ -types requires the naming of n parameters. (This example also illustrates why VC density is, arguably, a better measurement of complexity than the number of parameters in the defining formulas for types.)

Example 3.15. Let $\mathcal{L} = \{U, S, E\}$ be the language consisting of a unary predicate U and a binary predicates S and E . Fix $n > 0$, and let $\mathbf{M}_n = (M_n, U, S, E)$ be the \mathcal{L} -structure with the following properties (see Figure 3.1):

1. E is an equivalence relation with countably many classes, each of size $2n + \binom{2n}{n}$;
2. U splits every E -class C so that $|C \cap U| = 2n$ and $|C \cap \neg U| = \binom{2n}{n}$;
3. S holds only between E -equivalent elements and only between elements of U and $\neg U$; also, for each E -class C and each n -element subset X of $C \cap U$ let there be a unique b_X in $C \cap \neg U$ so that $S(b_X, y)$ holds if and only if $y \in X$, and let these be the only pairs in S .

The \mathcal{L} -theory $T_n = \text{Th}(\mathbf{M}_n)$ is totally categorical with $\text{MR}(T_n) = 1$, and Corollary 3.12 applies. For a fixed E -class C , if $Y \subseteq C \cap U$ has size less than n , then for any $a_1, a_2 \in (C \cap U) \setminus Y$ there is an automorphism of \mathbf{M}_n fixing Y and switching a_1 and a_2 . Thus fixing an E -class C and letting $A = C \cap U$, if Z is an A -definable set definable with fewer than n parameters from A , then $|Z \cap A| \neq n$. Hence for any $X \subseteq A$ the type $\text{tp}_{S(x,y)}(b_X/A)$ is not definable using fewer than n parameters from A , yet the VC density of $S(x, y)$ is 1 by Corollary 3.12 and $R_{S(x,y)}(x = x) = 1$.

4 Modules of Finite Breadth

Let R be a ring (with 1). In this paper, “ R -module” always means “left R -module.” Let \mathcal{L}_R be the language of R -modules and M be an R -module, construed as an

\mathcal{L}_R -structure in the natural way. Recall that an \mathcal{L}_R -formula $\varphi(x)$ is said to be positive-primitive (p.p., for short) if φ is equivalent (in the theory of R -modules) to a formula of the form $\exists y Ax = By$ where $A \in R^{k \times |x|}$, $B \in R^{k \times |y|}$, for some k (where we employed the usual shorthand notation of combining a conjunction of linear equations into a single matrix equation). By the Baur-Monk Theorem, each \mathcal{L}_R -formula is equivalent, in $\text{Th}(M)$, to a Boolean combination of p.p. \mathcal{L}_R -formulas. (Our main references for the model theory of modules are [33] and [17, Appendix A.1].) In this section we first introduce a notion of breadth of M , defined via its lattice of p.p. definable subgroups, and, then show that if M has finite breadth d , then M has the VC d property. (Our usage of breadth is not obviously related to the concept of breadth of a module considered in [33, Chapter 10].) We also show that uniform bounds on the VC density (in M) of formulas with m parameter variables can be computed from a certain quotient of the lattice of p.p. definable subgroups of M^m .

4.1 Dimension, width, height and breadth. We begin by recalling a few definitions and basic facts from the theory of (partially) ordered sets. In the following (P, \leq) denotes an ordered set. (We also simply denote (P, \leq) by P if the ordering \leq is understood from the context.) Given another ordered set (P', \leq') , a map $f: P \rightarrow P'$ is said to be increasing (or a morphism of ordered sets) if $a \leq b \Rightarrow f(a) \leq' f(b)$ for all $a, b \in P$ and strictly increasing if $a < b \Rightarrow f(a) < f(b)$ for all $a, b \in P$. We call f an embedding of ordered sets if $a \leq b \iff f(a) \leq' f(b)$ for all $a, b \in P$. Unless otherwise stated, proofs of the facts mentioned in the rest of this subsection can be found in any standard text on ordered sets (e.g., [16]).

The ordering \leq can be extended to a linear ordering on P (Szpilrajn/Marczewski). Moreover, if $\{\leq_i\}_{i \in I}$ is the family of all linear orderings on P extending \leq , then $\leq = \bigcap_{i \in I} \leq_i$, i.e., for all $a, b \in P$ one has $a \leq b$ iff $a \leq_i b$ for every $i \in I$. The *dimension* of P is the smallest non-zero cardinal $d = \dim(P)$ such that there is a family $\{\leq_i\}_{i \in I}$ of $d = |I|$ linear orderings on P with $\leq = \bigcap_{i \in I} \leq_i$. Equivalently, $\dim(P)$ is the smallest cardinal $d > 0$ such that P can be embedded into a direct product of d chains. Here and below, given a non-empty family $\{(P_i, \leq_i)\}_{i \in I}$ of ordered sets (P_i, \leq_i) , the direct product $\prod_{i \in I} P_i$ is equipped with the ordering defined by $(a_i) \leq (b_i)$ iff $a_i \leq_i b_i$ for every $i \in I$. We are mostly interested in finite-dimensional ordered sets. For all finite-dimensional ordered sets P and P' ,

$$\dim(P \times P') \leq \dim(P) + \dim(P'),$$

with equality if both P and P' have a greatest and a smallest element [41, Chapter 2, (3.4) and (3.5)].

The *width* of P is defined to be the supremum of the cardinalities of antichains in P , and denoted by $\text{width}(P)$. Clearly if P is non-empty, then P is a chain iff $\dim(P) = 1$ iff $\text{width}(P) = 1$. Dilworth's Theorem states that if $\text{width}(P)$ is finite, then it equals the smallest number of chains in P whose union is P . As a consequence of this theorem one obtains the inequality

$$\dim(P) \leq \text{width}(P), \tag{4.1}$$

provided P is of finite width. Dually, the *height* of P is defined to be the supremum of the cardinalities of a chain in P , denoted by $\text{height}(P)$.

The *breadth* of P , denoted by $\text{breadth}(P)$, is the smallest integer $d \geq 0$ (if it exists) with the following property: for all $x_1, \dots, x_{d+1}, y_1, \dots, y_{d+1} \in P$ such that $x_i \leq y_j$ for all $i \neq j$ in $[d+1]$ there exists $i \in [d+1]$ such that $x_i \leq y_i$. If there is no such d , then we set $\text{breadth}(P) = \infty$. This definition of breadth for an arbitrary ordered set (which is “self-dual”: $\text{breadth}(P, \leq) = \text{breadth}(P, \geq)$) is due to Wehrung. It is easy to see that if (L, \wedge) is a (meet-) semilattice, then the breadth of L may also be described as the smallest d (if it exists) such that for all $x_1, \dots, x_n \in L$ with $n > d$ there are $i_1, \dots, i_d \in [n]$ with $x_1 \wedge \dots \wedge x_n = x_{i_1} \wedge \dots \wedge x_{i_d}$. (Compare with Section 2.3.) Let P' be another ordered set; then

$$\text{breadth}(P \times P') \leq \text{breadth}(P) + \text{breadth}(P'),$$

with equality if P and P' have a greatest and a smallest element, and if P' embeds into P , then $\text{breadth}(P') \leq \text{breadth}(P)$. Together with (4.1) this yields the following fundamental inequality: if P has finite width, then

$$\text{breadth}(P) \leq \dim(P) \leq \text{width}(P). \quad (4.2)$$

Note that in general there is no bound on $\text{breadth}(P)$ in terms of $\text{height}(P)$: the “standard example St_d ” of a d -dimensional ordered set (see, e.g., [41, Chapter 1, (5.1)]) actually has breadth d , yet has height 2. However, if (L, \wedge) is a semilattice, then

$$\text{breadth}(L) \leq \text{height}(L), \quad (4.3)$$

with a strict inequality if L has a largest element. (See also [1, Example 2.16].)

We have $\text{breadth}(L') \leq \text{breadth}(L)$ if there exists a surjective morphism of semilattices $L \rightarrow L'$. A map $f: L \rightarrow L'$ between semilattices is a morphism of semilattices if $f(a \wedge b) = f(a) \wedge f(b)$ for all $a, b \in L$. (This condition is stronger than merely being a morphism of ordered sets, when L, L' are viewed only as ordered sets.) An injective morphism of semilattices is called an embedding of semilattices. By [6, Section II.5, Exercise 6 (c)] we have $\text{breadth}(L) \geq n$ iff L embeds the semilattice $(2^{[n]}, \cap)$.

4.2 Breadth in modular and distributive lattices. In this subsection, L denotes a lattice with smallest element 0 and largest element 1, and we assume $0 \neq 1$. Recall that L is said to be modular if the identity

$$(a \wedge c) \vee (b \wedge c) = ((a \wedge c) \vee b) \wedge c$$

holds for all $a, b, c \in L$. So far we have discussed upper bounds on the breadth of ordered sets and semilattices. Lower bounds on the breadth of a modular lattice can be obtained via its (dual) Goldie dimension, and for distributive lattices, the breadth can be computed as the width of an associated (often simpler) ordered set.

A subset A of $L \setminus \{0\}$ is said to be *join-independent* if $(\bigvee A') \wedge a = 0$ for all finite subsets A' of A and all $a \in A \setminus A'$. The *Goldie dimension* of L is the largest $n = \text{Gdim}(L)$ such that L contains a join-independent subset of size n , if there is such an n ; otherwise we set $\text{Gdim}(L) = \infty$. The Goldie dimension of the dual L^* of L is called the *dual Goldie dimension* of L and denoted by $\text{Gdim}^*(L)$. (See [15].)

Lemma 4.1. *Suppose L is modular. Then*

$$\text{breadth}(L) \geq \max \{ \text{Gdim}(L), \text{Gdim}^*(L) \}.$$

Proof Since the lattices L and L^* have the same breadth, it suffices to prove that $\text{breadth}(L) \geq \text{Gdim}(L)$. Let $A \subseteq L \setminus \{0\}$ be a finite join-independent set. For each $a \in A$ put $\hat{a} := \bigvee(A \setminus \{a\})$. Then for each $A' \subseteq A$ we have

$$\bigwedge_{a \in A'} \hat{a} = \bigwedge_{a \in A'} \bigvee(A \setminus \{a\}) = \bigvee \left(\bigcap_{a \in A'} A \setminus \{a\} \right) = \bigvee A \setminus A',$$

using [14, Chapter IV, Theorem 11]. In particular $\bigwedge_{a \in A} \hat{a} = 0 \neq \bigwedge_{a \in A'} \hat{a}$ for all proper subsets A' of A . This shows $\text{breadth}(L) \geq |A|$, so yields the claim. \square

Distributive lattices form an important subclass of the class of modular lattices. Recall that L is said to be distributive if

$$a \vee (b \wedge c) = (a \vee b) \wedge (a \vee c) \quad \text{for all } a, b, c \in L.$$

Equivalently, a lattice is distributive iff it is isomorphic to a sublattice of 2^X (with the usual meet and join), for some set X (Birkhoff-Stone Theorem). For finite distributive lattices, such an isomorphism can be described explicitly: Recall that a non-zero element a of L is join-irreducible if whenever $a = x \vee y$ with $x, y \in L$, then $a = x$ or $a = y$, and let $J(L)$ be the set of join-irreducible elements of L , equipped with the ordering induced by L . Also, for an ordered set P we denote by $I(P)$ the set of initial segments of P , ordered by inclusion, and for $a \in L$ let

$$(a) := \{x \in J(L) : x \leq a\} \in I(J(L)).$$

If L is finite and distributive, then the map $a \mapsto (a) : L \rightarrow I(J(L))$ is an isomorphism.

Proposition 4.2 (Dilworth, cf. [41, Chapter 2, (8.7)]). *Suppose L is finite and distributive. Then*

$$\text{breadth}(L) = \dim(L) = \text{width}(J(L)), \quad \text{height}(L) = |J(L)| + 1,$$

and $\text{Gdim}(L)$ is the maximum number of pairwise disjoint initial segments of $J(L)$.

Example. Suppose $L = 2^{[n]}$. Then L is distributive, and

$$\text{breadth}(L) = \dim(L) = \text{Gdim}(L) = n, \quad \text{height}(L) = n + 1.$$

It is also well-known that $\text{width}(L) = \binom{n}{\lfloor n/2 \rfloor}$ (Sperner's Theorem).

4.3 The breadth of the lattice of p.p. definable subgroups. Let $\text{PP} = \text{PP}_R$ be the set of p.p. \mathcal{L}_R -formulas in the single indeterminate x modulo equivalence in the theory of R -modules. In the following we don't distinguish between a p.p. \mathcal{L}_R -formula and its representative in PP . The set PP has the structure of a modular lattice with meet and join given by

$$\begin{aligned} (\varphi \wedge \psi)(x) &:= \varphi(x) \wedge \psi(x), \\ (\varphi \vee \psi)(x) &:= \exists y \exists z (\varphi(y) \wedge \psi(z) \wedge x = y + z). \end{aligned}$$

having a smallest element (represented by $x = 0$) and a largest element (represented by $x = x$). Given an R -module M , we also equip the collection of p.p. definable subgroups

$$\text{PP}_m(M) = \{\varphi(M) : \varphi(x) \text{ p.p. } \mathcal{L}_R\text{-formula with } |x| = m\}$$

of M^m with the structure of a modular lattice via

$$M_1 \wedge M_2 = M_1 \cap M_2, \quad M_1 \vee M_2 = M_1 + M_2 \quad \text{for } M_1, M_2 \in \text{PP}_m(M),$$

having smallest element $\{0\}$ and greatest element M^m . For all m_1, m_2 we have an embedding of lattices

$$\text{PP}_{m_1}(M) \times \text{PP}_{m_2}(M) \rightarrow \text{PP}_{m_1+m_2}(M) \tag{4.4}$$

given by

$$(H_1, H_2) \mapsto (H_1 \times \{0\}^{m_2}) \oplus (\{0\}^{m_1} \times H_2).$$

If N is a submodule of M and N is p.p. definable in M , then clearly $\text{PP}_m(N)$ is a sublattice of $\text{PP}_m(M)$, and denoting the natural surjection $M^m \rightarrow (M/N)^m$ by π , we have an injective lattice morphism

$$H \mapsto \pi^{-1}(H): \text{PP}_m(M/N) \rightarrow \text{PP}_m(M).$$

If N is a pure submodule of M (e.g., if N is a direct summand of M), then for every p.p. definable subgroup H of M^m , the subgroup $H \cap N^m$ is p.p. definable (by the same p.p. formula), and the map

$$H \mapsto H \cap N^m: \text{PP}_m(M) \rightarrow \text{PP}_m(N) \tag{4.5}$$

is a surjective morphism of lattices. If M is an R -module and I is a non-empty index set, then after choosing some arbitrary $i_0 \in I$ and identifying M with the i_0 -component of $N = M^{(I)}$ in the natural way, M becomes a direct summand of $N = M^{(I)}$, and the resulting map (4.5) is an isomorphism of lattices.

The following is also easy to verify (cf. [33, Lemma 2.10]):

Lemma 4.3. *Suppose $\{M_i\}_{i \in I}$ is a family of R -modules, where $I \neq \emptyset$. Then $M = \bigoplus_{i \in I} M_i$ is pure in $\prod_{i \in I} M_i$, and the map*

$$H \mapsto H \cap M^m: \text{PP}_m(\prod_{i \in I} M_i) \rightarrow \text{PP}_m(M)$$

is bijective. Moreover, for each $H \in \text{PP}_m(M)$ we have $H = \bigoplus_{i \in I} H \cap M_i$, and the map

$$H \mapsto (H \cap M_i)_{i \in I}: \text{PP}_m(M) \rightarrow \prod_{i \in I} \text{PP}_m(M_i) \tag{4.6}$$

is an embedding of lattices. If I is finite and each M_i is p.p. definable in M , then (4.6) is onto.

In the rest of this subsection we focus on $\text{PP}(M) = \text{PP}_1(M)$. The map $\varphi \mapsto \varphi(M): \text{PP} \rightarrow \text{PP}(M)$ is a surjective morphism of lattices. Note that there is an R -module M^* such that this map $\text{PP} \rightarrow \text{PP}(M^*)$ is an isomorphism of lattices (e.g., take $M^* =$ the direct sum of a set of representatives for all isomorphism types of finitely presented R -modules, cf. [33, Corollary 8.17]). In particular

$$\text{breadth}(\text{PP}_R) = \max_M \text{breadth}(\text{PP}(M)).$$

We say that M has breadth d if $\text{PP}(M)$ has breadth d , and similarly for width and height. (So PP_R has finite breadth iff there is a uniform [finite] bound on the breadth of all R -modules.) We collect some basic properties of breadth (immediate from the preceding discussions):

Lemma 4.4. *Let M, N be R -modules. Then:*

1. If N is a p.p. definable submodule of M , then

$$\text{breadth}(N), \text{breadth}(M/N) \leq \text{breadth}(M);$$

2. if N is a pure submodule of M , then $\text{breadth}(N) \leq \text{breadth}(M)$;
3. if I is a non-empty index set, then $\text{breadth}(M^{(I)}) = \text{breadth}(M)$;
4. $\text{breadth}(M \oplus N) \leq \text{breadth}(M) + \text{breadth}(N)$, with equality if both M and N are p.p. definable in $M \oplus N$.

It can be checked that the Goldie dimension $\text{Gdim}(\text{PP}(M))$ of the modular lattice $\text{PP}(M)$ is the largest n (if it exists) such that M contains an internal direct sum $H_1 \oplus \dots \oplus H_n$ of n non-zero p.p. definable subgroups H_1, \dots, H_n of M . The dual Goldie dimension $\text{Gdim}^*(\text{PP}(M))$ of $\text{PP}(M)$ is the largest n (if it exists) such that M contains proper p.p. definable subgroups H_1, \dots, H_n of M such that $H_i + \bigcap_{j \neq i} H_j = M$ for all $i \in [n]$. We refer to $\text{Gdim}(\text{PP}(M))$ and $\text{Gdim}^*(\text{PP}(M))$ as the p.p. Goldie dimension and the dual p.p. Goldie dimension of M , respectively. By Lemma 4.1 we have

$$\text{breadth}(M) \geq \max \{ \text{Gdim}(\text{PP}(M)), \text{Gdim}^*(\text{PP}(M)) \}.$$

The following is the archetypical example of a module of infinite breadth:

Example. Suppose $R = \mathbb{Z}$ and $M = \mathbb{Z}$, considered as a \mathbb{Z} -module as usual. Then $\text{Gdim}^*(\text{PP}(M)) = \infty$, since for each prime p the subgroup $p\mathbb{Z}$ is p.p. definable, and for all pairwise distinct primes p_1, \dots, p_d , we have $p_i\mathbb{Z} + \bigcap_{j \neq i} p_j\mathbb{Z} = \mathbb{Z}$ for all $i \in [d]$. Hence $\text{breadth}(M) = \infty$ and so $\text{breadth}(\text{PP}_{\mathbb{Z}}) = \infty$. (But note that $\text{Gdim}(\text{PP}(M)) = 1$.)

If M has only finitely many p.p. definable subgroups (e.g., if M is \aleph_0 -categorical), then clearly $\text{breadth}(M)$ is finite (and bounded by $\log_2|\text{PP}(M)|$). In fact, it is enough to require that M has finite height: $\text{breadth}(M) < \text{height}(M)$ by (4.3) (and this is strict since $\text{PP}(M)$ has a largest element). Note also that if $\text{height}(M)$ is finite, then $\text{MR}(T) < \text{height}(M)$, and if $M \equiv M^{\aleph_0}$ and $\text{MR}(T) < \omega$, then we have $\text{height}(M) = \text{MR}(T) + 1$ and hence $\text{breadth}(M) \leq \text{MR}(T)$; here $T = \text{Th}(M)$. (See [1, Example 3.14].)

Example. Prest [32] showed that all R -modules have finite Morley rank iff R is of finite representation type, i.e., there are finitely many indecomposable R -modules N_1, \dots, N_r such that every R -module is a direct sum of copies of the N_i . In this case there is a uniform (finite) bound d on the Morley rank of R -modules, hence $\text{breadth}(\text{PP}_R) \leq d$.

The previous example raises the following question, for which we do not know the answer:

Question 4.5. For which rings R do all R -modules have finite breadth, and for which rings R does PP_R have finite breadth?

If $\text{PP}(M)$ is finite and distributive, then Proposition 4.2 allows us to compute the breadth of $\text{PP}(M)$ as the width of the (often much simpler) ordered set $J(\text{PP}(M))$ of join-irreducibles of $\text{PP}(M)$. It is well-known (cf. [11, Theorem 3.1]) that if R is commutative, then for every R -module M the lattice $\text{PP}(M)$ is distributive iff R is a Prüfer ring, i.e., the lattice of all of its ideals is distributive. (Equivalently, a commutative ring R is Prüfer iff the ideals of each localization $R_{\mathfrak{m}}$ of R at a maximal ideal \mathfrak{m} of R are totally ordered by inclusion.)

4.4 Breadth and the VC d property. Let M be an R -module and $T = \text{Th}(M)$. The following observation, connecting the breadth of M with the VC d property, explains our interest in the concept of breadth:

Proposition 4.6. *Suppose M has (finite) breadth d . Then T has the VC d property. In particular, if in addition M is infinite, then $\text{vc}^T(m) \leq dm$ for every m .*

Proof Let Φ be the set of all p.p. \mathcal{L}_R -formulas in the single object variable x . By Baur-Monk, every \mathcal{L}_R -formula $\varphi(x; y)$ is equivalent in T to a Boolean combination of formulas from Φ . So if M has breadth d , then by Lemma 2.3, the set system \mathcal{S}_Φ has breadth d , hence by Lemma 2.5, every finite subset of Φ has UDTFS with d parameters; therefore T has the VC d property by Lemma 2.7. \square

Remark. A suitable modification of the argument in the proof above (using the natural multi-sorted version of Theorem 2.6, cf. [1, Corollary 5.8]) shows more generally that if \mathbf{G} is a 1-based expansion of a group, and the meet-semilattice of $\text{acl}^{\text{eq}}(\emptyset)$ -definable subgroups of G has breadth d , then (\mathbf{G}, G) has the VC d property, and hence $\text{vc}^{\text{Th}(\mathbf{G})}(m) \leq dm$ for every m .

We record two immediate corollaries of the previous proposition. The R -module M is called *p.p.-uniserial* if the ordered set $\text{PP}(M)$ is a chain.

Corollary 4.7. *If M is p.p.-uniserial, then T has the VC 1 property.*

The following corollary offers a more precise result than Corollary 3.14 (but with an identical bound on vc^T), under the additional assumption that $M^{\aleph_0} \equiv M$.

Corollary 4.8. *Suppose $M^{\aleph_0} \equiv M$. If M has finite Morley rank d , then M has the VC d property; in particular, if $\text{MR}(T) = 1$, then T has the VC 1 property.*

4.5 Commensurability and VC density. In the previous subsection, for the theory $T = \text{Th}(M)$ of the infinite R -module M we obtained an upper bound on the VC density $\text{vc}^T(m)$ in terms of m and the breadth of the lattice $\text{PP}(M) = \text{PP}_1(M)$. In this subsection we show that $\text{vc}^T(m)$ can indeed be computed precisely as the breadth of a suitable quotient of the lattice $\text{PP}_m(M)$. In fact, everything generalizes to 1-based groups, so we work in this wider setting. (In the last corollary of this subsection we additionally assume commutativity.)

Let G be a group (written multiplicatively). For subgroups H_1, H_2 of G we write $H_1 \lesssim H_2$ if $H_1 \cap H_2$ has finite index in H_1 . It is easy to see that \lesssim is a quasi-ordering on the set of subgroups of G which extends the ordering by inclusion. The equivalence relation \sim associated to \lesssim is called commensurability: $H_1 \sim H_2$ iff $H_1 \cap H_2$ is of finite index in both H_1 and H_2 . The \sim -class of a subgroup of G is called its commensurability class; the quasi-ordering \lesssim induces an ordering on the set of commensurability classes of subgroups of G . One easily verifies:

Lemma 4.9. *Let H_1 and H_2 be subgroups of G with $H_1 \lesssim H_2$. Then*

1. $H_1 \cap H \lesssim H_2 \cap H$ for every subgroup H of G ;
2. $H_1 H \lesssim H_2 H$ for every normal subgroup H of G .

In particular, by (1), \cap induces a semilattice structure on the set of commensurability classes of subgroups of G ; by (1) and (2), there is also a natural lattice structure on the set of commensurability classes of normal subgroups of G .

If G is finite, there is only one commensurability class; from now on assume G is infinite. Let also \mathbf{G} be an expansion of G (viewed as a structure in the language of groups as usual). *Throughout the rest of this subsection we assume $T = \text{Th}(\mathbf{G})$ is 1-based.* We then have:

Proposition 4.10. *Suppose that the semilattice of commensurability classes of $\text{acl}^{\text{eq}}(\emptyset)$ -definable subgroups of G^m has finite breadth d . Then $\text{vc}^T(m) \leq d$.*

Proof By Lemma 2.1 we may assume that the language of \mathbf{G} includes a constant symbol for every element of G . Let $\Delta(x; y)$ be a finite set of partitioned formulas where x is a tuple of variables of length m . By 1-basedness (see [19]) and Lemma 2.2 we may assume that each instance of a formula $\varphi \in \Delta$ defines a coset of an $\text{acl}^{\text{eq}}(\emptyset)$ -definable subgroup of G^m that we denote by H_φ . (Here and in the rest of the proof, “coset” means “left coset.”) Also for every $\Psi \subseteq \Delta$ we set $H_\Psi := \bigcap_{\varphi \in \Psi} H_\varphi$. We may also assume $|\Delta| \geq d$.

For each non-empty subset Ψ of Δ we choose $\Psi^\# \subseteq \Psi$ with $|\Psi^\#| \leq d$ such that $H_{\Psi^\#} \sim H_\Psi$. Let $\varphi_{\Psi,1}, \dots, \varphi_{\Psi,d}$ be the elements of $\Psi^\#$. Also, choose representatives $c_{\Psi,1}, \dots, c_{\Psi,n_\Psi} \in H_{\Psi^\#}$ for the cosets of H_Ψ in $H_{\Psi^\#}$. Set $n := \max_{\Psi \subseteq \Delta} n_\Psi$ where $n_\emptyset = 1$. For each Ψ and a coset E of $H_{\Psi^\#}$ in G^m we also choose an arbitrary element $c_\Psi(E)$ of E .

Let $B \subseteq G^{|y|}$ be finite and non-empty. We define a map

$$q \mapsto (\Psi_q, i_q, b_q): S^\Delta(B) \rightarrow 2^\Delta \times [n] \times B^d \quad (4.7)$$

as follows: given $q \in S^\Delta(B)$ let

$$\Psi = \Psi_q := \{\psi \in \Delta : \psi(x; b) \in q \text{ for some } b \in B\},$$

so

$$D = D_q := \bigcap_{\psi(x; b) \in q} \psi(G; b)$$

is a coset of H_Ψ . If $\Psi = \emptyset$ we let $i_q := 1$ and $b_q \in B^d$ be arbitrary. Suppose $\Psi \neq \emptyset$, and take $b = b_q = (b_1, \dots, b_d) \in B^d$ such that $\varphi_{\Psi,j}(x; b_j) \in q$ for $j = 1, \dots, d$. So D is contained in the coset $E = \bigcap_{j=1}^d \varphi_{\Psi,j}(G; b_j)$ of $H_{\Psi^\#}$, hence $c_\Psi(E)^{-1}D$ is a coset of H_Ψ contained in $H_{\Psi^\#}$ and thus of the form $c_\Psi(E)^{-1}D = c_{\Psi,i}H_\Psi$ for a unique $i = i_q \in [n_\Psi]$. We have

$$D = c_\Psi(E)c_{\Psi,i}H_\Psi \quad \text{where } E = \varphi_{\Psi,1}(x; b_1) \cap \dots \cap \varphi_{\Psi,d}(x; b_d),$$

so the triple (Ψ, i, b) uniquely determines q . Hence our map (4.7) is one-to-one. This shows that $\text{vc}^*(\Delta) \leq d$ as required. \square

It would be desirable if, strengthening Proposition 4.6, one could deduce linear growth of vc^T by simply knowing that the semilattice of commensurability classes of $\text{acl}^{\text{eq}}(\emptyset)$ -definable subgroups of G has finite breadth. An answer to the following natural question (which we do not know) would yield this:

Question 4.11. *Suppose the semilattice of commensurability classes of $\text{acl}^{\text{eq}}(\emptyset)$ -definable subgroups of G has finite breadth d . Does the structure \mathbf{G} have the VC d property? (Of course, the choice of the coset element $c_\Psi(E)$ in the proof above is the main obstacle.)*

This proposition gives rise to a characterization of dp-minimal 1-based groups (complementing the stability-theoretic characterization of stable dp-minimal theories given in [28]):

Corollary 4.12. *The following are equivalent:*

1. *T is dp-minimal;*
2. *T is vc-minimal: $\text{vc}^T(1) < 2$;*
3. $\text{vc}^T(1) = 1$;
4. *the set of commensurability classes of $\text{acl}^{\text{eq}}(\emptyset)$ -definable subgroups of G is linearly ordered by \lesssim .*

Proof The implication (4) \Rightarrow (3) holds by Proposition 4.10, (3) \Rightarrow (2) is trivial, and (2) \Rightarrow (1) holds by [1, Corollary 5.13], so it only remains to show (1) \Rightarrow (4). Suppose H_1, H_2 are $\text{acl}^{\text{eq}}(\emptyset)$ -definable subgroups of G which are incomparable with respect to \lesssim , i.e., $[H_i : H_1 \cap H_2] = \infty$ for each $i = 1, 2$. Let $H = H_1 \cap H_2$ and choose elements $b_{ij} \in H_{3-i} \setminus H$ ($i = 1, 2, j \in \mathbb{N}$) with $b_{ij}b_{ik}^{-1} \notin H$ for $i = 1, 2$ and distinct $j, k \in \mathbb{N}$. For $j_1, j_2 \in \mathbb{N}$ set $a_{j_1, j_2} = b_{1j_1} \cdot b_{2j_2}$. We also let $\varphi_1(x; y)$ be an \mathcal{L}^{eq} -formula stating that $y^{-1}x \in H_1$ and $\varphi_2(x; y)$ be an \mathcal{L}^{eq} -formula stating that $xy^{-1} \in H_2$. Then

$$\begin{aligned} \mathbf{G} \models \varphi_1(a_{j_1, j_2}; b_{1j_1}) &\iff b_{1j_1}^{-1}b_{1j_1}b_{2j_2} \in H_1 \\ &\iff b_{1j_1}^{-1}b_{1j_1} \in H_1 \\ &\iff j = j_1, \end{aligned}$$

and similarly $\mathbf{G} \models \varphi_2(a_{j_1, j_2}; b_{2j_2})$ iff $j = j_2$. Hence $\varphi_1(x, y), \varphi_2(x, y)$ and the sequences $(b_{ij})_{j \in \mathbb{N}}$ ($i = 1, 2$) form an ICT pattern in (\mathbf{G}, G) . This gives rise to an ICT pattern in \mathbf{G} , showing that \mathbf{G} is not dp-minimal. \square

We also obtain the description of $\text{vc}^T(m)$ advertised at the beginning of this subsection.

Corollary 4.13. *Suppose G is abelian. Then $\text{vc}^T(m)$ equals the breadth of the lattice of commensurability classes of $\text{acl}^{\text{eq}}(\emptyset)$ -definable subgroups of G^m .*

This is an immediate consequence of Proposition 4.10 and the following observation (in which we revert back to additive notation) generalizing the argument in the proof of (1) \Rightarrow (4) in Corollary 4.12:

Lemma 4.14. *Let A be an expansion of an infinite abelian group in a language \mathcal{L} expanding the language $\{0, +\}$ of abelian groups. Suppose A^m contains definable subgroups H_1, \dots, H_d , where $d > 1$, such that for each $i \in [d]$, the definable subgroup $H := H_1 \cap \dots \cap H_d$ has infinite index in*

$$H_{\neq i} := H_1 \cap \dots \cap H_{i-1} \cap H_{i+1} \cap \dots \cap H_d.$$

Then $\text{vc}^{\text{Th}(A)}(m) \geq d$.

Proof Let H_1, \dots, H_d be as in the hypothesis, and let i and k range over $[d]$. Let y be a tuple of new variables with $|x| = |y|$, and for each i let $\varphi_i(x; y)$ be an \mathcal{L} -formula stating that $x - y \in H_i$, and set $\Delta := \{\varphi_1, \dots, \varphi_d\}$. Since $[H_{\neq i} : H]$ is infinite, we may choose elements b_{ij} ($j \in \mathbb{N}$) of $H_{\neq i} \setminus H$ which are pairwise inequivalent modulo H . Fix $t \in \mathbb{N}$, and set

$$B := \{b_{ij} : i \in [d], j \in [t]\},$$

so $|B| = dt$. For each $\mathbf{j} = (j(1), \dots, j(d)) \in [t]^d$ set

$$a_{\mathbf{j}} = b_{1j(1)} + \dots + b_{dj(d)}.$$

Then for each i we have

$$a_{\mathbf{j}} - b_{ij(i)} = \sum_{k \neq i} b_{kj(k)} \in \sum_{k \neq i} H_{\neq k} \subseteq H_i.$$

(Here we use that A is abelian.) Hence for each i, k and $l \in [t]$ we have

$$\begin{aligned} A \models \varphi_i(a_{\mathbf{j}}; b_{kl}) &\iff a_{\mathbf{j}} - b_{kl} \in H_i \\ &\iff b_{ij(i)} - b_{kl} \in H_i \\ &\iff k = i \text{ and } l = j(i). \end{aligned}$$

So the $a_{\mathbf{j}}$ realize distinct $\Delta(x; B)$ -types in A . Thus $\pi_{\Delta}^*(dt) \geq |S^{\Delta}(B)| \geq t^d$, and since t was arbitrary, this yields $\text{vc}^{\text{Th}(A)}(m) \geq d$. \square

4.6 Commensurability and VC density for modules. We now return to the setting of modules. Let M be an infinite R -module and $T = \text{Th}(M)$. We write $\widetilde{\text{PP}}_m(M)$ for the lattice of commensurability classes of p.p. definable subgroups of M^m . We have a natural surjective lattice morphism $\text{PP}_m(M) \rightarrow \widetilde{\text{PP}}_m(M)$, so $\widetilde{\text{PP}}_m(M)$ is modular (and distributive if $\text{PP}_m(M)$ is distributive) with

$$\text{breadth}(\text{PP}_m(M)) \geq \text{breadth}(\widetilde{\text{PP}}_m(M)).$$

For all m_1, m_2 the natural embedding (4.4) shows that

$$\text{breadth}(\text{PP}_{m_1}(M)) + \text{breadth}(\text{PP}_{m_2}(M)) \leq \text{breadth}(\text{PP}_{m_1+m_2}(M))$$

and

$$\text{breadth}(\widetilde{\text{PP}}_{m_1}(M)) + \text{breadth}(\widetilde{\text{PP}}_{m_2}(M)) \leq \text{breadth}(\widetilde{\text{PP}}_{m_1+m_2}(M)). \quad (4.8)$$

If $M \equiv M^{\aleph_0}$, then the morphism $\text{PP}_m(M) \rightarrow \widetilde{\text{PP}}_m(M)$ is bijective (see, e.g., [17, Lemma A.1.7]) and hence $\text{PP}_m(M)$ and $\widetilde{\text{PP}}_m(M)$ have the same breadth, for each m .

Corollary 4.15. *For each m ,*

$$m \text{breadth}(\widetilde{\text{PP}}(M)) \leq \text{breadth}(\widetilde{\text{PP}}_m(M)) = \text{vc}^T(m) \leq m \text{breadth}(\text{PP}(M)),$$

with both inequalities being equalities if $M \equiv M^{\aleph_0}$.

Proof The first inequality is a consequence of (4.8), the equation is a special case of Corollary 4.13, and the last inequality is Proposition 4.6. \square

Question 4.16. *Is $\text{breadth}(\widetilde{\text{PP}}_m(M)) = m \text{breadth}(\widetilde{\text{PP}}(M))$ for every m ? (In this case we would have $\text{vc}^T(m) = m \text{vc}^T(1)$ for each m , in line with the Question posed in Section 3.2 of [1].)*

Example. Let K be an infinite field, $R = Ke_1 \oplus \dots \oplus Ke_d$ the ring-theoretic direct product of $d > 0$ copies of K , with idempotents e_1, \dots, e_d , and $M = R$ as a module over itself. Then M is the direct sum of its definable submodules

$$Ke_i = \{a \in M : e_j a = 0 \text{ for all } j \neq i\} \quad (i = 1, \dots, d),$$

so $\text{PP}(M) \cong 2^{[d]}$ by Lemma 4.3 and hence $\text{breadth}(M) = d$. In particular, M has the $\text{VC } d$ property by Proposition 4.6. Thus $\text{vc}^T(m) \leq dm$ for every m , and since $M \equiv M^{\aleph_0}$, we have in fact $\text{vc}^T(m) = dm$ for every m , by the previous corollary.

Since $\text{breadth}(\widetilde{\text{PP}}(M))$ bounds both the Goldie dimension $\text{Gdim}(\widetilde{\text{PP}}(M))$ and the dual Goldie dimension $\text{Gdim}^*(\widetilde{\text{PP}}(M))$ of the lattice $\widetilde{\text{PP}}(M)$, the previous corollary implies:

Corollary 4.17. *Let H_1, \dots, H_d be p.p. definable subgroups of M . If*

1. *H_i is infinite and $H_i \cap \sum_{j \neq i} H_j$ is finite, for each i , or*
2. *H_i has infinite index in M and $H_i + \bigcap_{j \neq i} H_j$ has finite index in M , for each i ,*

then $\text{vc}^T(m) \geq md$ for every m .

One also verifies easily the following properties of the breadth of the lattice of p.p. definable subgroups modulo commensurability (cf. Section 4.3):

Lemma 4.18. *Let M, N be R -modules. Then for each m :*

1. *If N is p.p. definable in M , then*

$$\text{breadth}(\widetilde{\text{PP}}_m(N)), \text{breadth}(\widetilde{\text{PP}}_m(M/N)) \leq \text{breadth}(\widetilde{\text{PP}}_m(M));$$

2. *if N is a pure submodule of M , then*

$$\text{breadth}(\widetilde{\text{PP}}_m(N)) \leq \text{breadth}(\widetilde{\text{PP}}_m(M));$$

3. *if I is a non-empty index set, then*

$$\text{breadth}(\widetilde{\text{PP}}_m(M^{(I)})) = \text{breadth}(\widetilde{\text{PP}}_m(M));$$

4. *and*

$$\text{breadth}(\widetilde{\text{PP}}_m(M \oplus N)) \leq \text{breadth}(\widetilde{\text{PP}}_m(M)) + \text{breadth}(\widetilde{\text{PP}}_m(N)),$$

with equality if both M and N are p.p. definable in $M \oplus N$, or if N is finite (in which case $\text{breadth}(\widetilde{\text{PP}}_m(N)) = 0$).

In Section 2.5 we noted that in general, dp-minimality is not preserved under interpretability. However, together with Corollary 4.12 the case $m = 1$ of Lemma 4.18 immediately yields that dp-minimality is preserved under some natural constructions with modules:

Corollary 4.19. *Let M be a dp-minimal R -module. Then the following R -modules are also dp-minimal:*

1. *any p.p. definable submodule N of M and its quotient M/N ;*
2. *any pure submodule of M ;*
3. *any power $M^{(I)}$, where I is a non-empty index set;*
4. *any direct sum $M \oplus N$ where N is a finite R -module; and*
5. *any direct summand of M of finite index.*

To finish this subsection we now relate the lattice $\widetilde{\text{PP}}_m(M)$ to another modular lattice (of connected \wedge -definable subgroups of M^m) which plays an important role in the model theory of modules. As usual we say that a subgroup H of M^m is (p.p. definably) *connected* if for each p.p. definable subgroup H' of M^m we have $H \lesssim H'$ iff $H \leq H'$ (i.e., $[H : H \cap H'] < \infty \Rightarrow H \subseteq H'$). In the following

we assume that M is $|R|^+$ -saturated. We let $\text{PP}_m^0(M)$ be the set of connected subgroups of M^m which arise as the intersection of a (possibly infinite) family of p.p. definable subgroups of M^m . (In [33] this is denoted by $\text{PP}_0^{(m)}(M)$.) We equip $\text{PP}_m^0(M)$ with the ordering by inclusion. One can show that then $\text{PP}_m^0(M)$ is a modular lattice; moreover, for $H \in \text{PP}_m(M)$ the subgroup

$$H^0 := \bigcap\{H' \in \text{PP}_m(M) : H \lesssim H'\}$$

of H is an element of $\text{PP}_m^0(M)$ (and called the connected component of H); see [33, Corollary 2.3 and Lemma 2.6]. The map

$$\text{PP}_m(M) \rightarrow \text{PP}_m^0(M) : H \mapsto H^0,$$

which is clearly a morphism of ordered sets, is in general not a morphism of lattices; in fact, the intersection of two connected p.p. definable subgroups of M is not in general connected [33, Example 2, p. 119]. Note however that two commensurable subgroups in $\text{PP}_m(M)$ have identical connected component, so $H \mapsto H^0$ factors through the canonical surjective morphism $\text{PP}_m(M) \rightarrow \widetilde{\text{PP}}_m(M)$:

$$\begin{array}{ccc} \text{PP}_m(M) & \xrightarrow{H \mapsto H^0} & \text{PP}_m^0(M) \\ \downarrow & \nearrow \iota & \\ \widetilde{\text{PP}}_m(M) & & \end{array}$$

The theory $T = \text{Th}(M)$ is superstable iff $\text{PP}_m^0(M) = \text{PP}_1^0(M)$ is well-founded, and in this case $\text{U}(T)$ is the foundation rank of $\text{PP}_m^0(M)$ [33, Corollary 5.13]; moreover:

Proposition 4.20. *Suppose T is superstable. Then $\widetilde{\text{PP}}_m(M)$ has the same foundation rank as $\text{PP}_m^0(M)$.*

If T is superstable we can employ the Shelah degree $R(-)$ (the ordinal-valued rank function denoted by $R^m(-, \mathcal{L}_R, \infty)$ in [38]), whose basic properties in this context we use freely (cf. [33, Section 5.2]). In particular we recall that for $G \in \text{PP}_m^0(M)$, the rank $R(G)$ is the foundation rank of G in $\text{PP}_m^0(M)$ [33, Lemma 5.10, Theorems 5.12 and 5.18]. The proposition is a consequence of the following two lemmas, in both of which we assume that T is superstable. The first lemma is immediate from [33, Lemma 5.10 and Corollary 5.15].

Lemma 4.21. *Let $G \in \text{PP}_m^0(M)$. Then there is some $K \in \text{PP}_m(M)$ with $G \leq K$ and $R(G) = R(K)$.*

Lemma 4.22. *The map ι is strictly increasing, i.e., for all $H, H' \in \text{PP}_m(M)$ we have $H \not\lesssim H' \Rightarrow H^0 \not\leq (H')^0$.*

Proof Let $H, H' \in \text{PP}_m(M)$; it is clear that $H \lesssim H' \Rightarrow H^0 \leq (H')^0$ (and this implication doesn't need superstability). So suppose for a contradiction that $H \not\lesssim H'$ and $H^0 = (H')^0$. Since $H \sim H' \cap H$, after replacing H by $H' \cap H$ we may assume $H \leq H'$. Then H has infinite index in H' , so $R(H) < R(H')$. Using Lemma 4.21, pick $K' \in \text{PP}_m(M)$ with $H' \not\lesssim K'$ and $R(K') = R((H')^0)$. Then

$$R(H) < R(H') = R(H' \cap K') \leq R(K') = R((H')^0) = R(H^0) \leq R(H),$$

a contradiction. \square

Proof of Proposition 4.20 By Lemma 4.22, the foundation rank of $\widetilde{\text{PP}}_m(M)$ is not larger than that of $\text{PP}_m^0(M)$. For the converse we show, by transfinite induction on the ordinal α , that if $G \in \text{PP}_m^0(M)$ satisfies $R(G) \geq \alpha$, then the foundation rank of each $K \in \text{PP}_m(M)$ with $K \geq G$ and $R(G) = R(K)$, viewed as element of $\widetilde{\text{PP}}_m(M)$, is no less than α . Since the case $\alpha = 0$ or α a limit ordinal are trivial, we only need to treat the case where $\alpha = \beta + 1$ is a successor ordinal. So let $G \in \text{PP}_m^0(M)$ with $R(G) \geq \alpha$ and let $K \in \text{PP}_m(M)$ with $K \geq G$ and $R(G) = R(K)$. Take $G' \in \text{PP}_m^0(M)$ with $G' \leq G$ and $R(G') = \beta$, and by Lemma 4.21 take $K' \in \text{PP}_m(M)$ with $G' \leq K'$ and $R(G') = R(K')$; after replacing K' by $K \cap K'$, we may assume $K' \leq K$. Now $R(K') < R(K)$ implies $[K : K'] = \infty$, hence the foundation rank of K in $\widetilde{\text{PP}}_m(M)$ is strictly larger than that of K' , and by inductive hypothesis the latter is greater than or equal to β . \square

Example 4.23. The module M has U-rank 0 iff M is finite, and M has U-rank 1 iff for each $H \in \text{PP}(M)$ we either have $H \sim 0$ or $H \sim M$.

Corollary 4.24. Suppose T has finite U-rank d . Then $\text{vc}^T(m) \leq dm$ for every m .

Proof By Proposition 4.20, $\widetilde{\text{PP}}_m(M)$ has height dm and hence breadth at most dm . Thus by Proposition 4.10 we get $\text{vc}^T(m) \leq dm$. \square

Remark. Every complete theory of a module is dimensional [33, Corollary 6.21] and hence does not have the finite cover property (see also [4]); hence Corollary 3.8 yields another proof of Corollary 4.24.

The relationship between the ordered sets $\widetilde{\text{PP}}_m(M)$ and $\text{PP}_m^0(M)$ is particularly clean if T is totally transcendental:

Lemma 4.25. If T is totally transcendental, then ι is an isomorphism of ordered sets.

Proof In this case, $\text{PP}_m(M)$ satisfies the descending chain condition [33, Theorem 3.1]; it follows that $H \sim H^0$ for all $H \in \text{PP}_m(M)$ (showing that ι is an embedding of ordered sets) and ι is onto (using [33, Lemma 5.10]). \square

4.7 Direct sums of modules. For use in the next section we establish two lemmas concerning the behavior of VC density under the operation of taking the direct sum of two modules. (We also prove a corresponding fact for the VC d property, which, however, will not be used later.) We let M, M' be R -modules, $T = \text{Th}(M)$, $T' = \text{Th}(M')$, and $T^+ = \text{Th}(M \oplus M')$. We first observe that if M is infinite and a pure submodule of M' , then $\text{vc}^T \leq \text{vc}^{T'}$. (This is an immediate consequence of the Baur-Monk Theorem and the fact that every Boolean combination of p.p. \mathcal{L}_R -formulas is invariant under the extension $M' \subseteq M$.) This yields:

Lemma 4.26. Suppose M is infinite and M' is finite. Then $\text{vc}^{T^+} = \text{vc}^T$.

Proof By the remark preceding the lemma, applied to $M \oplus M'$ in place of M' , we obtain $\text{vc}^{T^+} \geq \text{vc}^T$. The other inequality $\text{vc}^{T^+} \leq \text{vc}^T$ follows from the remark after Lemma 3.21 in [1]. \square

If both M and M' are infinite, by [1, Lemma 3.21] we can only conclude that

$$\max(\text{vc}^T, \text{vc}^{T'}) \leq \text{vc}^{T^+} \leq \text{vc}^T + \text{vc}^{T'}.$$

The standing assumption in both lemmas below is that M and M' are definable in $M \oplus M'$. Then the natural projections $\pi: M \oplus M' \rightarrow M$ and $\pi': M \oplus M' \rightarrow M'$ are also definable in $M \oplus M'$. In this situation, the estimate on vc^{T^+} above can be improved:

Lemma 4.27. *Suppose M and M' are infinite. Then $\text{vc}^{T^+} = \text{vc}^T + \text{vc}^{T'}$.*

Proof By Lemma 2.1 we may assume that M, M' are \emptyset -definable in $M \oplus M'$. We only need to show $\text{vc}^{T^+} \geq \text{vc}^T + \text{vc}^{T'}$. Let $\Delta(x; y), \Delta'(x; y)$ be finite sets of p.p. \mathcal{L}_R -formulas. Set (slightly abusing syntax)

$$\Delta_1 := \{\varphi(\pi(x); y) : \varphi(x; y) \in \Delta\}, \quad \Delta'_1 := \{\varphi'(\pi'(x); y) : \varphi'(x; y) \in \Delta'\}.$$

Let $B \subseteq M^{|y|}, B' \subseteq (M')^{|y|}$ with $|B| = |B'| = t \in \mathbb{N}$, and put

$$C = B \cup B' \subseteq (M \oplus M')^{|y|},$$

so $|C| \leq 2t$. Given a pair $(q, q') \in S^\Delta(B) \times S^{\Delta'}(B')$ we let p be the $(\Delta_1 \cup \Delta'_1)(x; C)$ -type $\text{tp}^{\Delta_1 \cup \Delta'_1}(a_q + a_{q'}/C)$ of $a_q + a_{q'} \in (M \oplus M')^{|x|}$ in $M \oplus M'$, where $a_q, a_{q'}$ are arbitrary realizations of q in M and of q' in M' , respectively. Then for each $\varphi \in \Delta$ and $b \in B$,

$$\begin{aligned} \varphi(x; b) \in q &\iff M \models \varphi(a_q; b) \\ &\iff M \oplus M' \models \varphi(a_q; b) \\ &\iff M \oplus M' \models \varphi(\pi(a_q + a_{q'}); b) \\ &\iff \varphi(\pi(x); b) \in p, \end{aligned}$$

where in the second equivalence we used that φ is p.p. Similarly we see that $\varphi'(x; b') \in q'$ iff $\varphi'(\pi'(x); b') \in p$, for each $\varphi'(x; y) \in \Delta'$ and $b' \in B'$. Hence the map

$$S^\Delta(B) \times S^{\Delta'}(B') \rightarrow S^{\Delta_1 \cup \Delta'_1}(C)$$

given by $(q, q') \mapsto p$ is injective, so $\pi_\Delta^*(t) \cdot \pi_{\Delta'}^*(t) \leq \pi_{\Delta_1 \cup \Delta'_1}^*(2t)$. Here π_Δ^* and $\pi_{\Delta'}^*$ are computed in M and M' , respectively, and $\pi_{\Delta_1 \cup \Delta'_1}^*$ is computed in $M \oplus M'$. Since t was arbitrary, this yields $\text{vc}^*(\Delta) + \text{vc}^*(\Delta') \leq \text{vc}^*(\Delta_1 \cup \Delta'_1)$, employing a similar convention for vc^* as for π^* . By Lemma 2.2 (and the Baur-Monk Theorem) we obtain $\text{vc}^T(m) + \text{vc}^{T'}(m) \leq \text{vc}^{T^+}(m)$, where $m = |x|$, as required. \square

Lemma 4.28. *Suppose M and M' are \emptyset -definable in $M \oplus M'$. If T has the VC d property and T' has the VC d' property, then T^+ has the VC $(d + d')$ property.*

Proof Suppose T has the VC d property and T' has the VC d' property. Let $\Delta(x; y)$, where $|x| = 1$, be a finite set of p.p. \mathcal{L}_R -formulas; by Lemma 2.7 and the Baur-Monk Theorem it suffices to show that Δ has UDTFS with $d + d'$ parameters in $M \oplus M'$. Let $\mathcal{D} = (\mathcal{D}_i)_{i \in I}$ and $\mathcal{D}' = (\mathcal{D}'_{i'})_{i' \in I'}$ be uniform definitions of $\Delta(x; B)$ -types over finite sets in M and in M' , respectively, where

$$\mathcal{D}_i := \{d_{\varphi, i}(y; \bar{y}) : \varphi \in \Delta\}, \quad \mathcal{D}'_{i'} := \{d'_{\varphi, i'}(y; \bar{y}) : \varphi \in \Delta\}$$

with $\bar{y} = (y_1, \dots, y_d)$ and $\bar{y}' = (y'_1, \dots, y'_{d'})$. By Baur-Monk we may assume that each $d_{\varphi, i}$ and $d'_{\varphi, i'}$ is a Boolean combination of p.p. \mathcal{L}_R -formulas.

Let C be a finite set of tuples from $(M \oplus M')^{|y|}$. Take $B \subseteq M^{|y|}$, $B' \subseteq (M')^{|y|}$ minimal such that each $c \in C$ is of the form $c = b + b'$ for a pair $(b, b') \in B \times B'$. Let $a \in M$, $a' \in M'$ and $p = \text{tp}^\Delta(a + a'/C)$. Put

$$q := \text{tp}^\Delta(a/B), \quad q' := \text{tp}^\Delta(a'/B'),$$

and take $\bar{b} = (b_1, \dots, b_d) \in B^d$, $\bar{b}' = (b'_1, \dots, b'_{d'}) \in (B')^{d'}$ and $i \in I$, $i' \in I'$ such that $\mathcal{D}_i(y; \bar{b})$ defines q and $\mathcal{D}'_{i'}(y; \bar{b}')$ defines q' . Then for all $(b, b') \in B \times B'$ we have

$$\begin{aligned} \varphi(x; b + b') \in p &\iff M \oplus M' \models \varphi(a + a'; b + b') \\ &\iff M \models \varphi(a; b) \text{ and } M' \models \varphi(a'; b') \\ &\iff M \models d_{\varphi, i}(b; \bar{b}) \text{ and } M' \models d'_{\varphi, i'}(b'; \bar{b}'), \end{aligned}$$

where in the second equivalence we used that φ is p.p. Take $\bar{c} \in C^d$ and $\bar{c}' \in C^{d'}$ such that $\pi(\bar{c}) = \bar{b}$ and $\pi'(\bar{c}') = \bar{b}'$. Since $d_{\varphi, i}$ and $d'_{\varphi, i'}$ are Boolean combinations of p.p. formulas, we have

$$M \models d_{\varphi, i}(b; \bar{b}) \iff M \oplus M' \models d_{\varphi, i}(\pi(b + b'); \pi(\bar{c}))$$

and

$$M' \models d'_{\varphi, i'}(b'; \bar{b}') \iff M \oplus M' \models d'_{\varphi, i'}(\pi'(b + b'); \pi'(\bar{c}')).$$

Hence $\mathcal{D}^+ = (\mathcal{D}_{(i, i')}^+)_{(i, i') \in I \times I'}$, where $\mathcal{D}_{(i, i')}^+ = \{d_{\varphi, (i, i')}^+ : \varphi \in \Delta\}$ with

$$d_{\varphi, (i, i')}^+(y; \bar{y}, \bar{y}') := d_{\varphi, i}(\pi(y); \pi(\bar{y})) \wedge d'_{\varphi, i'}(\pi'(y); \pi'(\bar{y})),$$

is a uniform definition of $\Delta(x; B)$ -types over finite sets in $M \oplus M'$ with $d + d'$ parameters, as required. \square

5 Abelian Groups with Uniformly Bounded VC Density

In this final section of the paper we focus on abelian groups. We write abelian groups additively and construe them as first-order structures in the language $\mathcal{L}_{\mathbb{Z}}$ of \mathbb{Z} -modules as usual. Throughout this section A denotes an abelian group and $T = \text{Th}(A)$ its complete theory. Determining the VC density function of an arbitrary (infinite) abelian group is an interesting but probably intricate problem. In the first part of this section we obtain a satisfactory answer for those abelian groups whose lattice of p.p. definable subgroups is finite; in particular, we prove Theorem 1.2 from the introduction. After a preliminary subsection (Section 5.2) we then prove Theorem 1.3 in Section 5.3, and in Section 5.4 we characterize all dp-minimal abelian groups. We finish with some remarks on (non-) dp-minimal expansions of the dp-minimal abelian group \mathbb{Z} (Section 5.5).

5.1 A case study: abelian groups of finite exponent. For each prime p let

$$U(p; A) := \{i \geq 0 : U(p, i; A) \neq 1\}.$$

Then A has finite exponent iff A is torsion and each set $U(p; A)$ is finite with $U(p; A) = \emptyset$ for all but finitely many p . We will show:

Theorem 5.1. *Suppose A has finite exponent, and set*

$$d := \sum_p d(U(p; A)). \tag{5.1}$$

Then A has breadth d .

Here, given a finite non-empty set $I = \{i_1, \dots, i_n\}$ of integers $i_1 < \dots < i_n$, we define $d = d(I)$ as the maximal length of a sequence $1 = j(1) < j(2) < \dots < j(d) \leq n$ of indices such that for $k = 1, \dots, d-1$,

$$j(k+1) - j(k) = \begin{cases} 1 & \text{if } i_{j(k)+1} - i_{j(k)} > 1, \\ 2 & \text{otherwise.} \end{cases}$$

We also write $d(i_1, \dots, i_n)$ instead of $d(I)$, and set $d(\emptyset) := 0$.

Examples 5.2.

1. For $I = \{2, 3, 5, 7, 8, 9\}$ the sequence of $j(k)$'s is 1, 3, 4, 6, so $d(I) = 4$.
2. For $I = \{1, 2, 3, \dots, n\}$ the sequence of $j(k)$'s is $1, 3, 5, \dots, n$ if n is odd, and $1, 3, 5, \dots, n-1$ if n is even, hence $d(I) = \lceil n/2 \rceil$.

Clearly we have $d(I) \geq \lceil n/2 \rceil$ for every finite set I of n positive integers. We also have the following upper bound on d (which, however, is strict in general):

Lemma 5.3. *Let $I = \{i_1, \dots, i_n\}$ be as above, with $i_1 > 0$, and set $i_0 := 0$. Then*

$$d(I) \leq \min_{0 \leq j \leq n} (n - j + \lceil i_j/2 \rceil).$$

Proof Set $d_k := d(i_1, \dots, i_k)$ for $k \in [n]$, and let

$$1 = j(1) < j(2) < \dots < j(d_n) \leq n$$

be a sequence of indices as above; note that necessarily $j(d_k) = k$ or $j(d_k) = k-1$, and $d_k = d_{k-1}$ or $d_k = d_{k-1} + 1$. We proceed by induction to show that for each $k \in [n]$,

$$d_k \leq \min_{0 \leq j \leq k} (k - j + \lceil i_j/2 \rceil). \quad (5.2)$$

For $k = 1$ this is trivial, since $d_1 = 1$ and the right-hand side in this inequality is $\min\{1, \lceil i_1/2 \rceil\} = 1$. Suppose we have shown (5.2) for some value of $k < n$; then

$$d_{k+1} \leq d_k + 1 \leq (k+1) - j + \lceil i_j/2 \rceil \quad \text{for } j = 0, \dots, k,$$

so we only need to show that $d_{k+1} \leq \lceil i_{k+1}/2 \rceil$. For this, we distinguish several cases. Suppose first that $j(d_k) = k$ and $i_{k+1} - i_k > 1$. Then $d_{k+1} = d_k + 1$ (and $j(d_{k+1}) = j(d_k) + 1$), and $\lceil i_{k+1}/2 \rceil - \lceil i_k/2 \rceil \geq 1$, so

$$d_{k+1} \leq d_k + 1 \leq \lceil i_k/2 \rceil + 1 \leq \lceil i_{k+1}/2 \rceil.$$

If $j(d_k) = k$ and $i_{k+1} - i_k = 1$, then clearly

$$d_{k+1} = d_k \leq \lceil i_k/2 \rceil \leq \lceil i_{k+1}/2 \rceil.$$

Finally, suppose $j(d_k) = k-1$; then necessarily $k > 1$ and $i_k - i_{k-1} = 1$, $d_k = d_{k-1}$. In this case we have $d_{k+1} = d_k + 1$, $j(d_{k+1}) = k+1$, and clearly $i_{k+1} - i_{k-1} \geq 2$, so

$$d_{k+1} = d_{k-1} + 1 \leq \lceil i_{k-1}/2 \rceil + 1 \leq \lceil i_{k+1}/2 \rceil$$

as required. \square

In particular we have

$$\lceil n/2 \rceil \leq d(i_1, \dots, i_n) \leq \min\{n, \lceil i_n/2 \rceil\}. \quad (5.3)$$

For later use also note the following easily verified observation:

$$d(i_1, \dots, i_n) = 1 \iff n = 1, \text{ or } n = 2 \text{ and } i_2 = i_1 + 1. \quad (5.4)$$

5.1.1 Towards the proof of Theorem 5.1. Recall that every abelian torsion group is an internal direct sum of its p -primary components, as p ranges over the set of prime numbers. We denote the p -primary component of A by $A_p \leq A$; if A_p has finite exponent, then A_p is obviously p.p. definable in A . We have

$$A_p \cong \left(\bigoplus_{i>0} \mathbb{Z}(p^i)^{(\alpha_{p,i-1})} \right) \oplus \mathbb{Z}(p^\infty)^{(\beta_p)}$$

where $\alpha_{p,i}$ and β_p are cardinals, with $U(p, i; A) = p^{\alpha_{p,i}}$. (Here and below, $\mathbb{Z}(p^i)$ denotes the cyclic group of order p^i , and $\mathbb{Z}(p^\infty)$ the Prüfer p -group.)

Suppose now that A has finite exponent. Then by Lemma 4.4, (4) we have

$$\text{breadth}(A) = \sum_p \text{breadth}(A_p),$$

and in the representation of A_p above we have $\alpha_{p,i} = 0$ for all but finitely many pairs (p, i) , and $\beta_p = 0$ for each p . So for the proof of Theorem 5.1 we may assume that for some prime p , our group A has the following form:

$$A = \mathbb{Z}(p)^{(\alpha_0)} \oplus \mathbb{Z}(p^2)^{(\alpha_1)} \oplus \cdots \oplus \mathbb{Z}(p^m)^{(\alpha_{m-1})} \quad \text{where the } \alpha_i \text{ are cardinals.}$$

Letting

$$B = \mathbb{Z}(p)^{(\beta_0)} \oplus \mathbb{Z}(p^2)^{(\beta_1)} \oplus \cdots \oplus \mathbb{Z}(p^m)^{(\beta_{m-1})} \quad \text{where } \beta_i = \begin{cases} 0 & \text{if } \alpha_i = 0 \\ 1 & \text{otherwise,} \end{cases}$$

for an arbitrary infinite cardinal $\kappa \geq \max_i \alpha_i$ we have $A^{(\kappa)} \cong B^{(\kappa)}$ and thus

$$\text{breadth}(A) = \text{breadth}(A^{(\kappa)}) = \text{breadth}(B^{(\kappa)}) = \text{breadth}(B)$$

by part (3) of Lemma 4.4. Also $U(p; A) = U(p; B)$. Thus, after replacing A by B we may further assume that for each i , either $\alpha_i = 0$ or $\alpha_i = 1$, so

$$A = \mathbb{Z}(p^{i_1}) \oplus \mathbb{Z}(p^{i_2}) \oplus \cdots \oplus \mathbb{Z}(p^{i_n})$$

where $i_1 < \cdots < i_n$ are the indices $i \in [m]$ such that $\alpha_{i-1} \neq 0$. Hence (switching notation from i_k to λ_k) we are done once we have shown the following:

Proposition 5.4. *Suppose*

$$A = \mathbb{Z}(p^{\lambda_1}) \oplus \mathbb{Z}(p^{\lambda_2}) \oplus \cdots \oplus \mathbb{Z}(p^{\lambda_n})$$

is a finite abelian p -group of type $\lambda = (\lambda_1, \dots, \lambda_n)$, where $0 < \lambda_1 < \cdots < \lambda_n$. Then $\text{breadth}(A) = d(\lambda)$.

From now on assume that A is as in this proposition. We first give an explicit description of the lattice of p.p. definable subgroups of A .

5.1.2 The lattice of p.p. definable subgroups of a finite abelian p -group. The set \mathbb{N}^n , equipped with the product ordering, is a distributive lattice. We say that an n -tuple $\mu = (\mu_1, \dots, \mu_n) \in \mathbb{N}^n$ is *slowly growing* (relative to λ) if

$$0 \leq \mu_i - \mu_{i-1} \leq \lambda_i - \lambda_{i-1} \quad \text{for } i \in [n], \text{ where we set } \mu_0 := \lambda_0 := 0.$$

Note that if μ is slowly growing, then $\mu_i \leq \lambda_i$ for each $i \in [n]$, and $\mu_j - \mu_i \leq \lambda_j - \lambda_i$ for $i, j \in [n]$ with $i < j$. One verifies easily that the set Λ of slowly growing tuples is a (finite) sublattice of \mathbb{N}^n with smallest element 0 and largest element λ . If μ is slowly growing, then so is $\mu^* := \lambda - \mu$, and the map $\mu \mapsto \mu^*$ is an order-reversing involution of Λ , so $\Lambda^* \cong \Lambda$.

For each integer $k \geq 0$ consider the n -tuple $\mu = (\mu_1, \dots, \mu_n) \in \mathbb{N}^n$, given by $\mu_i = \lambda_i$ if $\lambda_i \leq k$ and $\mu_i = k$ otherwise; we denote this n -tuple μ by $\langle k \rangle$. The map $k \mapsto \langle k \rangle$ is an embedding of the chain $K := \{0, \dots, \lambda_n\}$ into Λ .

Lemma 5.5. *The lattice Λ is generated by the elements $\langle k \rangle$, $\langle k \rangle^*$ where k ranges over K .*

Proof Let Λ' be the sublattice of Λ generated by the elements $\langle k \rangle$, $\langle k \rangle^*$ where $k \in K$. Let an arbitrary $\mu \in \Lambda$ be given. For each $i \in [n]$ set

$$\mu_{\leq i} = (\mu_1, \mu_2, \dots, \mu_i, \mu_i, \dots, \mu_i) \in \Lambda.$$

We show by induction on i that $\mu_{\leq i} \in \Lambda'$ for each $i \in [n]$. For $i = 1$ this is clear since $\mu_{\leq 1} = (\mu_1, \dots, \mu_1) = \langle \mu_1 \rangle$, and for the inductive step notice that if $i \in [n-1]$, then

$$\mu_{\leq i+1} = \langle \mu_{i+1} \rangle \wedge (\mu_{\leq i} \vee \langle \lambda_{i+1} - \mu_{i+1} \rangle^*).$$

In particular $\mu = \mu_{\leq n} \in \Lambda'$. \square

We write the elements of A as n -tuples $a = (a_1, \dots, a_n)$ where each a_i is an element of $\mathbb{Z}(p^{\lambda_i})$. For each $\mu \in \mathbb{N}^n$ now define the subgroup

$$A_\mu := \{a = (a_1, \dots, a_n) \in A : p^{\mu_i} a_i = 0 \text{ for each } i \in [n]\}$$

of A . (So, e.g., $A_0 = \{0\}$ and $A_\lambda = A$.) It is clear that $A_{\mu \wedge \nu} = A_\mu \cap A_\nu$ and $A_{\mu \vee \nu} = A_\mu + A_\nu$ for all $\mu, \nu \in \mathbb{N}^n$, i.e., $\mu \mapsto A_\mu$ is a morphism of lattices.

Lemma 5.6. *The map $\mu \mapsto A_\mu$ restricts to a lattice isomorphism $\Lambda \rightarrow \text{PP}(A)$.*

Proof For each $k \in K$ the subgroup $A_{\langle k \rangle}$ of A is defined by the p.p. formula $p^k x = 0$, and $A_{\langle k \rangle^*}$ is defined by $p^k|x$. Hence by the previous lemma, the image of the restriction of $\mu \mapsto A_\mu$ to Λ is contained in $\text{PP}(A)$. Conversely, it is well-known (and easy to see) that the lattice PP_R , where R is the ring $\mathbb{Z}/p^{\lambda_n}\mathbb{Z} = \mathbb{Z}(p^{\lambda_n})$, is generated by the p.p. formulas having the form $p^k x = 0$ or $p^k|x$ where $k \in K$. Thus our morphism maps onto $\text{PP}(A)$, and it is also clearly one-to-one. \square

5.1.3 The ordered set of join-irreducibles. Now that we have identified $\text{PP}(A)$ as Λ , we turn to giving an explicit description of the ordered set $J(\Lambda)$ of join-irreducibles of the distributive lattice Λ . (We are interested in such a description since $\text{breadth}(L) = \text{width}(J(L))$, as we recall from Proposition 4.2.) Consider

$$P := \{(i, j) : j \in [n], i \in [\lambda_j]\},$$

equipped with the ordering given by

$$(i, j) \leq (i', j') \iff i \leq i' \text{ and } \lambda_j - i \geq \lambda_{j'} - i'.$$

See Figure 5.1 for an example. Note that $|P| = \lambda_1 + \dots + \lambda_n$, and P has smallest element $(1, n)$ and largest element (λ_n, n) .

For each $(i, j) \in P$ we define $\mu = \mu(i, j) \in \mathbb{N}^n$ by

$$\mu_k := \begin{cases} \lambda_k - (\lambda_j - i) & \text{if } k < j, \\ i & \text{if } k \geq j. \end{cases}$$

Here $a \dot{-} b = \max\{a - b, 0\}$ for integers a, b . It is easy to verify that for each $(i, j) \in P$, the tuple $\mu(i, j)$ is slowly growing. Figure 5.2 illustrates $\mu(4, 4)$ in the example from Figure 5.1.




Figure 5.1 An example for (the underlying set of) the ordered set P

Lemma 5.7. *The map $(i, j) \mapsto \mu(i, j): P \rightarrow \Lambda$ is an embedding of ordered sets.*

Proof Let $(i, j), (i', j') \in P$; we need to show that $(i, j) \leq (i', j')$ if and only if $\mu(i, j) \leq \mu(i', j')$. Write $\mu(i, j) = (\mu_1, \dots, \mu_n)$ and $\mu(i', j') = (\mu'_1, \dots, \mu'_n)$. It is convenient to distinguish two cases. First assume that $j \geq j'$. Then we have $(i, j) \leq (i', j')$ iff $i \leq i'$. If $\mu(i, j) \leq \mu(i', j')$, then $i = \mu_j \leq \mu'_{j'} = i'$. Conversely, suppose $i \leq i'$. Then for each $k \in [n]$, we have

$$\begin{cases} \mu_k = i \leq i' = \mu'_k & \text{if } k \geq j, \\ \mu_k = \lambda_k - (\lambda_j - i) < i \leq i' = \mu'_k & \text{if } j' \leq k < j, \\ \mu_k = \lambda_k - (\lambda_j - i) \leq \lambda_k - (\lambda_{j'} - i') = \mu'_k & \text{if } k < j'. \end{cases}$$

This shows $\mu(i, j) \leq \mu(i', j')$. On the other hand, if we assume that $j < j'$, then $(i, j) \leq (i', j')$ iff $\lambda_j - i \geq \lambda_{j'} - i'$, and in a similar way as in the previous case one sees that this is equivalent to $\mu(i, j) \leq \mu(i', j')$. \square

Lemma 5.8. *Let $\mu = \mu(i, j)$ where $(i, j) \in P$, and $\nu \in \Lambda$ with $\nu < \mu$. Then $\nu_j < \mu_j$.*

Proof Suppose for a contradiction that $\nu_j = \mu_j$. Then $\nu_k = \mu_k = i$ for all $k \geq j$, so if we let $k_1 \in [n]$ be the largest index such that $\nu_{k_1} < \mu_{k_1}$, then $k_1 < j$. Let also $k_0 \in [n]$ be minimal such that $\lambda_{k_0} > \lambda_j - i$; note that $k_0 \leq j$. Since $\nu_k = \mu_k = 0$ for $k < k_0$, we have $k_0 \leq k_1 < j$. Then

$$\nu_{k_1+1} - \nu_{k_1} \leq \lambda_{k_1+1} - \lambda_{k_1} = \mu_{k_1+1} + \lambda_j - i - (\mu_{k_1} + \lambda_j - i) = \mu_{k_1+1} - \mu_{k_1}$$

and hence

$$\nu_{k_1+1} \leq \mu_{k_1+1} + \nu_{k_1} - \mu_{k_1} < \mu_{k_1+1},$$

contradicting the maximality of k_1 . \square

Let $\mu \in \Lambda$. Call $j \in [n]$ *critical* (for μ) if $\mu_j > \mu_{j-1}$ and either $j = n$ or $j < n$ and $\mu_{j+1} - \mu_j < \lambda_{j+1} - \lambda_j$. Clearly if j is critical, then $\mu - e_j \in \Lambda$. (Here and below, e_1, \dots, e_n denote the standard basis vectors in \mathbb{R}^n .) Moreover, one easily shows:




Figure 5.2 An example for $\mu(4,4)$

Lemma 5.9. Let $\mu \in \Lambda$ and $j, j' \in [n]$ with $j' < j$, and suppose all indices $j', j' + 1, \dots, j$ are non-critical.

1. If $j = n$ or $j < n$ and $\mu_j = \mu_{j+1}$, then $\mu_{j'-1} = \mu_{j'} = \dots = \mu_j$.
2. If $\mu_{k-1} < \mu_k$ for $k = j', \dots, j$ and $j < n$, then $\mu_k = \lambda_k - (\lambda_{j+1} - \mu_{j+1})$ for $k = j', \dots, j$.

Note that part (1) of the previous lemma shows in particular that μ has no critical index iff $\mu = 0$. On the other hand, if μ is join-irreducible, then μ cannot have more than one critical index; for if $j \neq j'$ both are critical, then $\mu = (\mu - e_j) \vee (\mu - e_{j'})$ where $\mu - e_j, \mu - e_{j'} \in \Lambda \setminus \{\mu\}$. Hence if μ is join-irreducible, then μ has exactly one critical index.

Lemma 5.10. Let $\mu \in \Lambda$ have the unique critical index $j \in [n]$; then $\mu = \mu(i, j)$ where $i = \mu_j$.

Proof By part (1) of the previous lemma we have $i = \mu_j = \mu_{j+1} = \dots = \mu_n$; note that $i > 0$ since $\mu \neq 0$. On the other hand, if $\mu_{k-1} = \mu_k$ for some $k \in [j]$, then by the same part of Lemma 5.9 we have $0 = \mu_1 = \dots = \mu_k$. So if we let $k_0 \in [n]$ be minimal such that $\mu_{k_0} > 0$, then $k_0 \leq j$ and $\mu_{k-1} < \mu_k$ for $k = k_0, \dots, j$ and hence $\mu_k = \lambda_k - (\lambda_j - \mu_j)$ for $k = k_0, \dots, j-1$. Also,

$$\lambda_{k_0} - (\lambda_j - \mu_j) = \mu_{k_0} - \mu_{k_0-1} \leq \lambda_{k_0} - \lambda_{k_0-1}$$

and hence $\lambda_{k_0-1} - (\lambda_j - \mu_j) \leq 0$. This yields the claim. \square

Corollary 5.11. $J(\Lambda) = \{\mu(i, j) : (i, j) \in P\}$.

Proof Lemma 5.8 implies that the $\mu = \mu(i, j)$ with $(i, j) \in P$ are join-irreducible: if $\nu \in \Lambda$ satisfies $\nu < \mu$, then $\nu \leq \mu - e_j$, in particular, there cannot exist $\nu, \nu' \in \Lambda$




Figure 5.3 Partition of P into the chains C_k

with $\mu = \nu \vee \nu'$ and $\nu, \nu' < \mu$. Conversely, by the previous lemma, every join-irreducible element μ of Λ is of the form $\mu = \mu(i, j)$ with $(i, j) \in P$. \square

5.1.4 The width of $J(\Lambda)$. Since by Lemma 5.7 and Corollary 5.11 we now have an explicit description of the ordered set $J(\Lambda)$, we can read off the various numerical invariants of Λ from P by means of Proposition 4.2. For example, we have

$$\text{height}(\Lambda) = \lambda_1 + \cdots + \lambda_n + 1 \quad \text{and} \quad \text{Gdim}(\Lambda) = \text{Gdim}^*(\Lambda) = 1.$$

To complete the proof of Proposition 5.4, it remains to compute the width of P (and hence of $J(\Lambda)$), using Dilworth's Theorem: we will specify a partition of P into $d = d(\lambda_1, \dots, \lambda_n)$ many chains which, by virtue of Dilworth's Theorem, shows $\text{width}(P) \leq d$, and, then we pick an element from each of these chains to form an antichain of size d , which shows $\text{width}(P) \geq d$.

We consider the strictly increasing sequence

$$1 = j(1) < \cdots < j(k) < \cdots < j(d)$$

defined according to the definition of $d = d(\lambda_1, \dots, \lambda_n)$, i.e., for $k \in [d-1]$, we have $j(k+1) = j(k) + 1$ if $\lambda_{j(k)+1} - \lambda_{j(k)} \geq 2$ and $j(k+1) = j(k) + 2$ otherwise. Now for $k \in [d]$ define the subsets

$$\begin{aligned} C_k^\downarrow &:= \{(k, j) : j = j(k), \dots, n\} \\ C_k^{\rightsquigarrow} &:= \{(i, j) : i = k, \dots, \lambda_{j(k)} - k + 1, j(k) \leq j < j(k+1)\} \\ C_k^{\nearrow} &:= \{(\lambda_j - k + 1, j) : j = j(k), \dots, n\} \end{aligned}$$

of \mathbb{N}^2 , where we set $j(d+1) = n+1$. Using Lemma 5.3 it is not difficult to verify that each of the sets C_k^\downarrow , C_k^{\rightsquigarrow} and C_k^{\nearrow} is contained in P . Note that in P , we have

$$(k, n) < (k, n-1) < \cdots < (k, j(k)),$$

so C_k^\downarrow is a chain in P . Also, C_k^{\rightsquigarrow} is a chain, since

$$(k, j(k)) < (k+1, j(k)) < \cdots < (\lambda_{j(k)} - k + 1, j(k)) \quad \text{if } j(k+1) - j(k) = 1$$

and

$$(k, j(k) + 1) < (k, j(k)) < (k + 1, j(k) + 1) < (k + 1, j(k)) < \dots < \\ (\lambda_{j(k)} - k + 1, j(k)) \quad \text{if } j(k + 1) - j(k) = 2.$$

Similarly we see that C_k^\nearrow is a chain, since

$$(\lambda_{j(k)} - k + 1, j(k)) < (\lambda_{j(k)+1} - k + 1, j(k) + 1) < \dots < (\lambda_n - k + 1, n).$$

Hence

$$C_k := C_k^\downarrow \cup C_k^\rightsquigarrow \cup C_k^\nearrow$$

is a chain, with smallest element (k, n) and largest element $(\lambda_n - k + 1, n)$. Figure 5.3 shows the C_k in the example introduced in Figure 5.1.

As with $d = d(\lambda_1, \dots, \lambda_n)$, the definition of the ordered set P and the chains C_k depend on $\lambda = (\lambda_1, \dots, \lambda_n)$; we write $P = P(\lambda)$ and $C_k = C_k(\lambda)$, respectively, to make this dependence explicit.

Lemma 5.12. *The chains C_1, \dots, C_d form a partition of P .*

Proof We proceed by induction on n . Note that

$$C_k \cap ([\lambda_n] \times \{n\}) = \begin{cases} \{k, \lambda_n - k + 1\} \times \{n\} & \text{if } j(k) < n - 1 \\ \{k, \dots, \lambda_n - k + 1\} \times \{n\} & \text{if } j(k) = n \text{ or } j(k) = n - 1, \end{cases}$$

and hence the sets $C_k \cap ([\lambda_n] \times \{n\})$ ($k = 1, \dots, d$) partition $[\lambda_n] \times \{n\}$. Suppose $n > 1$ and set $\lambda' = (\lambda_1, \dots, \lambda_{n-1})$; then we clearly have

$$P(\lambda') = P(\lambda) \cap ([\lambda_{n-1}] \times [n - 1])$$

as ordered sets, and

$$C_k(\lambda) \cap ([\lambda_{n-1}] \times [n - 1]) = \begin{cases} C_k(\lambda') & \text{if } j(k) \leq n - 1 \\ \emptyset & \text{if } j(k) = n. \end{cases}$$

By inductive hypothesis therefore, if $j(d) < n$, then the sets $C_k \cap ([\lambda_{n-1}] \times [n - 1])$ ($k = 1, \dots, d$) partition $P \cap ([\lambda_{n-1}] \times [n - 1])$, if $j(d) = n$, then the sets $C_k \cap ([\lambda_{n-1}] \times [n - 1])$ ($k = 1, \dots, d - 1$) partition $P \cap ([\lambda_{n-1}] \times [n - 1])$, with $C_d \subseteq [\lambda_n] \times \{n\}$. Hence the C_k ($k = 1, \dots, d$) partition $P = \bigcup_{k \in [n]} [\lambda_k] \times \{k\}$. \square

By Lemma 5.12 and Dilworth's Theorem we have $\text{width}(P) \leq d$; the next lemma thus shows that $\text{width}(P) = d$. (In Figure 5.3 the elements of this antichain are circled.)

Lemma 5.13. *The elements $(1, j(1)), \dots, (d, j(d))$ form an antichain of P .*

Proof For all $k \in [d - 1]$ we have $\lambda_{j(k+1)} - \lambda_{j(k)} \geq 2$, hence the sequence $\lambda_{j(k)} - k$ ($k \in [d]$) is strictly increasing. Since $(k, j(k)) \leq (l, j(l))$ iff $k \leq l$ and $\lambda_{j(k)} - k \geq \lambda_{j(l)} - l$ for all $k, l \in [d]$, this immediately yields the claim. \square

This finishes the proof of Proposition 5.4, and hence of Theorem 5.1. \square

From Theorem 5.1 and Proposition 4.6 we see immediately:

Corollary 5.14. *Suppose A has finite exponent, and let d be as in (5.1); then A has the VC_d property. Hence if A is infinite, then $\text{vc}^T(m) \leq dm$ for each m .*

For the following examples also consult Examples 5.2 above:

Example. Suppose

$$A = \mathbb{Z}(p^2)^{(\alpha_2)} \oplus \mathbb{Z}(p^3)^{(\alpha_3)} \oplus \mathbb{Z}(p^5)^{(\alpha_5)} \oplus \mathbb{Z}(p^7)^{(\alpha_7)} \oplus \mathbb{Z}(p^8)^{(\alpha_8)} \oplus \mathbb{Z}(p^9)^{(\alpha_9)},$$

where the α_i are non-zero cardinals. Then A has breadth 4 and hence A has the VC 4 property; so if one of the α_i is infinite, then $\text{vc}^T(m) \leq 4m$ for each m .

Example. Suppose

$$A = \mathbb{Z}(p)^{(\alpha_1)} \oplus \mathbb{Z}(p^2)^{(\alpha_2)} \oplus \cdots \oplus \mathbb{Z}(p^n)^{(\alpha_n)} \quad (n > 0)$$

where the α_i are non-zero cardinals. Then A has the VC d property where $d = \lceil n/2 \rceil$; so if A is infinite, then $\text{vc}^T(m) \leq \lceil n/2 \rceil m$ for each m .

In fact, we can now precisely determine the VC density function of an infinite abelian group of finite exponent. Recall that for each p , in the introduction we defined

$$U_{\geq \aleph_0}(p; A) = \{i \geq 0 : U(p, i; A) \geq \aleph_0\},$$

a subset of $U(p; A)$. With this notation we have:

Corollary 5.15. *Suppose A is \aleph_0 -categorical, and let*

$$d_0 := \sum_p d(U_{\geq \aleph_0}(p; A)).$$

Then $\text{vc}^T(m) = d_0 m$ for each m .

Proof We may write $A = A_{\geq \aleph_0} \oplus B$ where $U_{\geq \aleph_0}(p; A) = U(p; A_{\geq \aleph_0})$ for each p and B is finite: if we decompose each p -primary component as

$$A_p \cong \bigoplus_{i>0} \mathbb{Z}(p^i)^{(\alpha_{p,i-1})},$$

then

$$A_{\geq \aleph_0} \cong \bigoplus_{\substack{p \text{ prime, } i > 0 \\ \text{with } \alpha_{p,i-1} \geq \aleph_0}} \mathbb{Z}(p^i)^{(\alpha_{p,i-1})}, \quad B \cong \bigoplus_{\substack{p \text{ prime, } i > 0 \\ \text{with } \alpha_{p,i-1} < \aleph_0}} \mathbb{Z}(p^i)^{(\alpha_{p,i-1})}.$$

Replacing A by $A_{\geq \aleph_0}$ if necessary we may assume that $U_{\geq \aleph_0}(p; A) = U(p; A)$ for every p and so $A^{\aleph_0} \cong A$ and $d_0 = d = \text{breadth}(A)$. Now $\text{vc}^T(m) = dm$ follows from Corollary 4.15. \square

Theorem 1.2 from the introduction is a consequence of Corollary 5.15 and the upper bound in (5.3). We also obtain:

Corollary 5.16. *An \aleph_0 -categorical abelian group is dp-minimal iff it is isomorphic to one of the form*

$$\mathbb{Z}(p^k)^{(\alpha)} \oplus \mathbb{Z}(p^{k+1})^{(\beta)} \oplus B,$$

where p is a prime, $k > 0$, α, β are cardinals, at least one of which is infinite, and B is a finite abelian group.

Proof Suppose A is \aleph_0 -categorical. Then by the previous corollary,

$$\text{vc}^T(1) = \sum_p d(U_{\geq \aleph_0}(p; A)).$$

Thus by Corollary 4.12, A is dp-minimal iff $U_{\geq \aleph_0}(p; A) \neq \emptyset$ for exactly one prime p and for this p , $d(U_{\geq \aleph_0}(p; A)) = 1$. By (5.4), the latter condition is equivalent to $U_{\geq \aleph_0}(p; A) = \{k\}$ or $U_{\geq \aleph_0}(p; A) = \{k, k+1\}$, for some $k > 0$. \square

5.2 Asymptotics of p.p. definable subgroups in homocyclic groups. In this subsection we go beyond finite-exponent groups and consider abelian groups of the form

$$A = \bigoplus_{i>0} \mathbb{Z}(p^i)^{(\alpha_{i-1})}$$

where the α_{i-1} are natural numbers; we assume that A is infinite, i.e., $\alpha_{i-1} > 0$ for infinitely many $i > 0$. Our goal is to show:

Proposition 5.17. *For every m we have $\text{breadth}(\widetilde{\text{PP}}_m(A)) \leq m$.*

Note that by Corollary 4.15 (and since always $\text{vc}^T(m) \geq m$), this proposition implies $\text{vc}^T(m) = m$ for every m . Before we give a proof of Proposition 5.17, we first look at the instructive case $m = 1$, in which we also have a sharper result. Below, for each n we let $A[n] := \{a \in A : na = 0\}$ (a p.p. definable subgroup of A).

Lemma 5.18. *The ordered set $\widetilde{\text{PP}}(A) = \widetilde{\text{PP}}_1(A)$ is a chain of order type $\omega + \omega^*$, with representatives given by the p.p. definable subgroups $A[p^d]$ and $p^e A$:*

$$0 \precsim A[p] \precsim \cdots \precsim A[p^d] \precsim A[p^{d+1}] \precsim \cdots \precsim p^{e+1} A \precsim p^e A \precsim \cdots \precsim pA \precsim A.$$

Proof For all $d < d'$ the p.p. formula $\delta_{d,d'}^p(x) = \delta_{d,d'}(x) := (p^{d'}|p^d x)$ defines the subgroup

$$\delta_{d,d'}(\mathbb{Z}(p^i)) = \begin{cases} \mathbb{Z}(p^i) & \text{if } i \leq d \\ p^{i-d}\mathbb{Z}(p^i) & \text{if } d < i < d' \\ p^{d'-d}\mathbb{Z}(p^i) & \text{if } i \geq d' \end{cases}$$

of $\mathbb{Z}(p^i)$, hence

$$\delta_{d,d'}(A) = \bigoplus_{i < d'} p^{\max\{0, i-d\}} \mathbb{Z}(p^i)^{(\alpha_{i-1})} \oplus \bigoplus_{i \geq d'} p^{d'-d} \mathbb{Z}(p^i)^{(\alpha_{i-1})}.$$

Note that setting $e := d' - d$, the group $\delta_{d,d'}(A)$ is commensurable with

$$p^e A = \bigoplus_{i \geq e} p^e \mathbb{Z}(p^i)^{(\alpha_{i-1})}.$$

For each d the p.p. formula $\tau_d^p(x) = \tau_d(x) := (p^d x = 0)$ defines the subgroup $\tau_d(\mathbb{Z}(p^i)) = p^{\max\{i-d, 0\}} \mathbb{Z}(p^i)$ in $\mathbb{Z}(p^i)$, so

$$A[p^d] = \tau_d(A) \sim \bigoplus_{i \geq d} p^{i-d} \mathbb{Z}(p^i)^{(\alpha_{i-1})}.$$

This description of the subgroups $A[p^d]$ and $p^e A$ makes it clear that $A[p^d] \precsim p^e A$ for all d, e . Since every p.p. definable subgroup of A is a finite intersection of groups definable by τ_d 's or $\delta_{d,d'}$'s, the lemma follows. \square

The key in this proof was that for any choice of d and e , for large enough i we have an inclusion $p^e \mathbb{Z}(p^i) \geq p^{i-d} \mathbb{Z}(p^i)$ among the subgroups of $\mathbb{Z}(p^i)$ defined by the p.p. formulas $p^e|x$ and $p^d x = 0$, respectively. For the case $m > 1$ we proceed in a similar way, by first investigating the intersection behavior, as $i \rightarrow \infty$, of subgroups of the homocyclic p -group $\mathbb{Z}(p^i)^m$ defined by a fixed collection of p.p. formulas:

Lemma 5.19. *Let $\varphi_1(x), \dots, \varphi_n(x)$ be p.p. $\mathcal{L}_{\mathbb{Z}}$ -formulas where $|x| = m$. There are $i_1, \dots, i_m \in [n]$ such that $\varphi_1 \wedge \dots \wedge \varphi_n$ and $\varphi_{i_1} \wedge \dots \wedge \varphi_{i_m}$ define the same subgroup of $\mathbb{Z}(p^i)^m$, for all sufficiently large i .*

In the proof of this lemma, which we give after some preliminary observations, we use the following basic properties of finitely generated modules over discrete valuation rings (DVRs):

Lemma 5.20. *Let R be a DVR with maximal ideal \mathfrak{m} and residue field $k = R/\mathfrak{m}$, let M be a finitely generated R -module, and let $\overline{M} = M/\mathfrak{m}M$, viewed as a k -linear space, with natural surjection $x \mapsto \overline{x}: M \rightarrow \overline{M}$.*

1. *Elements a_1, \dots, a_n of M form a minimal generating set for the R -module M iff $\overline{a_1}, \dots, \overline{a_n}$ form a basis for the k -linear space \overline{M} .*
2. *If M is torsion-free, then M is free, and any minimal generating set for M is a basis for M .*
3. *If M is a submodule of R^m , then M can be generated by m elements.*

In particular, if M is a submodule of R^m and $a_1, \dots, a_n \in M$ generate M , then there are $i_1, \dots, i_r \in [n]$, where $r = \dim_k \overline{M} \leq m$, such that a_{i_1}, \dots, a_{i_r} is a basis for M .

We omit the proof of these facts, which can be found in any standard text on commutative algebra. (Part (1) is a consequence of Nakayama's Lemma and holds more generally if R is a local ring.) We use this lemma to show:

Lemma 5.21. *Let R be a DVR and π a generator of the maximal ideal of R , and let $a_1, \dots, a_n, b_1, \dots, b_n \in R^m$. There are $i_1, \dots, i_r, j_1, \dots, j_s \in [n]$ with $r + s \leq m$ and some integer $e \geq 0$ such that*

1. *a_{i_1}, \dots, a_{i_r} generate the R -submodule $N = Ra_1 + \dots + Ra_n$ of R^m ;*
2. *for each $i \in [n]$ we have $\pi^e b_i \in N + R\pi^e b_{j_1} + \dots + R\pi^e b_{j_s}$.*

Proof By the lemma above, first choose $i_1, \dots, i_r \in [n]$ such that a_{i_1}, \dots, a_{i_r} is a basis for N . For each k let M_k be the R -submodule $N + R\pi^k b_1 + \dots + R\pi^k b_n$ of R^m , so

$$N \subseteq \dots \subseteq M_{k+1} \subseteq M_k \subseteq \dots \subseteq M_0 =: M.$$

Now the torsion submodule of M/N is finitely generated, hence we may take some $e \in \mathbb{N}$ such that if $a \in M$ satisfies $\pi^k a \in N$ for some $k \geq 0$, then $\pi^e a \in N$. Then M_e/N is torsion-free: if $a \in M_e$ is torsion in M_e/N , then $\pi^e a \in N$, so writing $a = b + \pi^e c$ ($b \in N$, $c \in M$) we have $\pi^{2e} c = \pi^e a - \pi^e b \in N$ and hence $\pi^e c \in N$, i.e., $a \in N$. Now apply the lemma again and choose $j_1, \dots, j_s \in [n]$ such that $\pi^e b_{j_1}, \dots, \pi^e b_{j_s}$ map onto a basis for M_e/N under the natural surjection $M_e \rightarrow M_e/N$. Then $M_e = \bigoplus_k Ra_{i_k} \oplus \bigoplus_l R\pi^e b_{j_l}$ (internal direct sum of R -submodules of M_e), so $r + s \leq m$ (again by the lemma above), and the i_k , j_l and e have the desired properties. \square

We now prove Lemma 5.19:

Proof of Lemma 5.19 Let $a_1, \dots, a_n, b_1, \dots, b_n \in \mathbb{Z}^m$ and $d_1, \dots, d_n \in \mathbb{N}$, and consider the p.p. formulas

$$\chi_j(x) = (a_j x = 0), \quad \psi_j(x) = (p^{d_j} | b_j x) \quad (j \in [n]).$$

Here, for $a = (a_1, \dots, a_m) \in \mathbb{Z}^m$ and $x = (x_1, \dots, x_m)$, we denote by ax the $\mathcal{L}_{\mathbb{Z}}$ -term $a_1x_1 + \dots + a_mx_m$. It suffices to show (cf. [17, Lemma A.2.1] or [33, Theorem 2.Z1]): there are $i_1, \dots, i_r, j_1, \dots, j_s \in [n]$ with $r+s \leq m$ such that for sufficiently large i , the formulas $\chi_1 \wedge \dots \wedge \chi_n \wedge \psi_1 \wedge \dots \wedge \psi_n$ and $\chi_{i_1} \wedge \dots \wedge \chi_{i_r} \wedge \psi_{j_1} \wedge \dots \wedge \psi_{j_s}$ define the same subgroup of $\mathbb{Z}(p^i)^m$. Now $\psi_j(x)$ and $p^{\max\{0, i-d_j\}}b_jx = 0$ define the same subgroup of $\mathbb{Z}(p^i)^m$, and if i and d are such that $i \geq d \geq d_j$ for each j , then $p^{\max\{0, i-d_j\}}b_j = p^{i-d}(p^{d-d_j}b_j)$, hence after replacing b_j by $p^{d-d_j}b_j$ we may assume that $d = d_j$ for each j . Set $R = \mathbb{Z}_{(p)}$ (the localization of the ring of integers at its prime ideal $(p) = p\mathbb{Z}$) and $\pi = p$, and choose $i_1, \dots, i_r, j_1, \dots, j_s \in [n]$ and $e \in \mathbb{N}$ as in Lemma 5.21. Then by condition (1) in this lemma, $\chi := \chi_1 \wedge \dots \wedge \chi_n$ and $\chi_{i_1} \wedge \dots \wedge \chi_{i_r}$ define the same subgroup of $\mathbb{Z}(p^i)^m$; similarly, by condition (2), if $i \geq d+e$, then $\chi \wedge \psi_1 \wedge \dots \wedge \psi_n$ and $\chi \wedge \psi_{j_1} \wedge \dots \wedge \psi_{j_s}$ define the same subgroup of $\mathbb{Z}(p^i)^m$. \square

Lemma 5.19 immediately implies Proposition 5.17. \square

We note that if the question posed after the proof of Proposition 4.10 had a positive answer, then Proposition 5.17 would simply follow from Lemma 5.18.

5.3 Abelian groups with uniform bounds on VC density. Throughout this subsection we assume that A is infinite. A list of abelian groups A with $\text{vc}^T(m) = m$ for all m (so that in particular, $T = \text{Th}(A)$ is dp-minimal) includes: \mathbb{Z}^n and $\mathbb{Z}^n \oplus \mathbb{Q}$, for each n (in each case, an expansion has the VC1 property, by [1, Corollary 6.5]); and any p.p. uniserial abelian group such as

$$\mathbb{Z}(p^d)^{(\alpha)} \oplus \mathbb{Z}_{(p)}^{(\beta)} \oplus \mathbb{Q}^{(\gamma)}$$

where d is a positive integer or ∞ , and α, β, γ are cardinals (by Corollary 4.7). (Here $\mathbb{Z}_{(p)}$ is the additive group of the localization of the ring of integers at its prime ideal $(p) = p\mathbb{Z}$.) On the other hand, the following lemma shows that for many A we have $\text{vc}^T(1) = \infty$.

Lemma 5.22. *Suppose p_1, \dots, p_d are pairwise distinct primes. If*

1. $A[p_i]$ is infinite for each i , or
2. A/p_iA is infinite for each i ,

then $\text{vc}^T(m) \geq dm$ for each m (and so if $d > 1$, then T is not dp-minimal).

Proof In the first case, this follows from Corollary 4.17, (1) applied to $H_i = A[p_i]$, and in the second case, from Corollary 4.17, (2) applied to $H_i = p_iA$. \square

In this subsection we extend this observation to a characterization of all A satisfying $\text{vc}^T(m) < \infty$ for all m . We say that a prime p is *non-singular* for A if both $A[p]$ and A/pA are finite, and *singular* for A otherwise. We also say that an abelian group is *non-singular* if it is either finite, or infinite and each prime is non-singular for it, and *singular* otherwise. For example, the additive group \mathbb{Z} of integers is non-singular. If an infinite abelian group is non-singular, then it has U-rank 1 (cf. Example 4.23); but of course, the converse of this implication does not hold, as witnessed, e.g., by infinite elementary abelian p -groups.

Theorem 5.23. *The following are equivalent:*

1. $\text{vc}^T(1) < \infty$;

2. $\text{vc}^T(m) < \infty$ for every m ;
3. there is some d such that $\text{vc}^T(m) \leq dm$ for every m ;
4. there are only finitely many p which are singular for A , and for all p the set $U_{\geq \aleph_0}(p; A)$ is finite.

For the proof of this theorem, recall that a *Szmielew group* is a countable abelian group of the form

$$\bigoplus_{p \text{ prime}} \left(\bigoplus_{n>0} \mathbb{Z}(p^n)^{(\alpha_{p,n-1})} \oplus \mathbb{Z}(p^\infty)^{(\beta_p)} \oplus \mathbb{Z}_{(p)}^{(\gamma_p)} \right) \oplus \mathbb{Q}^{(\delta)}$$

where the $\alpha_{p,n-1}$, β_p , γ_p and δ are cardinals (finite or \aleph_0). Such a Szmielew group is *strict* if

1. δ is either 0 or \aleph_0 ;
2. if $\beta_p \neq 0$ or $\gamma_p \neq 0$ for some p or $\alpha_{p,n-1} \neq 0$ for infinitely many pairs (p, n) , then $\delta = 0$; and
3. for each p , if there is no finite upper bound on the order of the elements which have p -power order and are not divisible by p , then $\beta_p = \gamma_p = 0$.

Any abelian group is elementarily equivalent to a unique strict Szmielew group (cf. [17, A.2]). Also, suppose B is an abelian group, elementarily equivalent to A , and p is a prime. Then for each $N \in \mathbb{N}$ we have $|A[p]| = N \iff |B[p]| = N$ and $|A/pA| = N \iff |B/pB| = N$, so p is non-singular for A iff p is non-singular for B . Similarly one sees that $U_{\geq \aleph_0}(p; A) = U_{\geq \aleph_0}(p; B)$. Hence A satisfies condition (4) in Theorem 5.23 iff B does. So for the proof of the theorem we may assume that A is a strict Szmielew group as above. We then have

$$U(p, n; A) = |(p^n A)[p]/(p^{n+1} A)[p]| = p^{\alpha_{p,n}} \quad \text{for all } n,$$

and

$$\dim_{\mathbb{F}_p} A[p] = \sum_{i \geq 0} \alpha_{p,i} + \beta_p, \quad \dim_{\mathbb{F}_p} A/pA = \sum_{i \geq 0} \alpha_{p,i} + \gamma_p.$$

We see from this that p is non-singular for A iff $\alpha_{p,n} < \aleph_0$ for every n , with $\alpha_{p,n} = 0$ for all but finitely many n , and $\beta_p, \gamma_p < \aleph_0$. Hence if $U_{\geq \aleph_0}(p; A)$ is non-empty, then p is singular for A .

Proof of Theorem 5.23. Since the implications (3) \Rightarrow (2) \Rightarrow (1) are trivial, we only need to prove (1) \Rightarrow (4) \Rightarrow (3). Suppose that $\text{vc}^T(1) < \infty$. By Lemma 5.22 there are only finitely many singular primes for A . Suppose p is such that $U(p, i; A) \geq \aleph_0$ for infinitely many i . Let $(i_k)_{k>0}$ be a strictly increasing sequence of non-negative integers such that $U(p, i_k; A) \geq \aleph_0$ for each k ; then

$$A_n := \bigoplus_{k \in [n]} \mathbb{Z}(p^{i_k})^{(\alpha_{p,i_k-1})}$$

is a pure subgroup of A , so $\text{vc}^T \geq \text{vc}^{T_n}$ for $T_n = \text{Th}(A_n)$; see the discussion at the beginning of Section 4.7. Since A_n has finite exponent p^{i_n} and $U_{\geq \aleph_0}(p; A_n) = \{i_1, \dots, i_n\}$, by Corollary 5.15 and (5.3) we have $\text{vc}^{T_n}(1) \geq \lceil n/2 \rceil$. Since this holds for each n , we obtain $\text{vc}^T(1) = \infty$, a contradiction. This shows (1) \Rightarrow (4).

For the proof of the remaining implication $(4) \Rightarrow (3)$, suppose condition (4) holds. Then the set

$$\mathcal{P} = \{p : \dim_{\mathbb{F}_p} A[p] < \aleph_0 \text{ and } \dim_{\mathbb{F}_p} A/pA < \aleph_0\}$$

of non-singular primes for A contains all but finitely many primes. The group

$$A^{[\mathcal{P}]} := \bigoplus_{p \in \mathcal{P}} \left(\bigoplus_{n>0} \mathbb{Z}(p^n)^{(\alpha_{p,n-1})} \oplus \mathbb{Z}(p^\infty)^{(\beta_p)} \oplus \mathbb{Z}_{(p)}^{(\gamma_p)} \right) \oplus \mathbb{Q}^{(\delta)}$$

is non-singular. Thus if $A^{[\mathcal{P}]}$ is infinite, then, with $T^{[\mathcal{P}]} = \text{Th}(A^{[\mathcal{P}]})$, we have

$$\text{vc}^{T^{[\mathcal{P}]}}(m) = m \quad \text{for every } m, \quad (5.5)$$

by Corollary 4.24. Now let $p \notin \mathcal{P}$ and set

$$A^{[p]} := \bigoplus_{n>0} \mathbb{Z}(p^n)^{(\alpha_{p,n-1})} \oplus \mathbb{Z}(p^\infty)^{(\beta_p)} \oplus \mathbb{Z}_{(p)}^{(\gamma_p)}$$

and $T^{[p]} = \text{Th}(A^{[p]})$, and suppose $A^{[p]}$ is infinite. If $\alpha_{p,n} = 0$ for all but finitely many n , then $A^{[p]}$ is the direct sum of a finite exponent group and the p.p. uniserial abelian group $\mathbb{Z}(p^\infty)^{(\beta_p)} \oplus \mathbb{Z}_{(p)}^{(\gamma_p)}$, and hence satisfies

$$\text{vc}^{T^{[p]}}(m) \leq m \cdot (d(U_{\geq \aleph_0}(p; A)) + 1) \quad \text{for each } m, \quad (5.6)$$

by Corollaries 5.15 and 4.7, respectively. Suppose $\alpha_{p,n} > 0$ for infinitely many n ; then $\beta_p = \gamma_p = 0$ since A is strict Szemielew, and by assumption (4) the set $U_{\geq \aleph_0}(p; A)$ is finite. Decomposing $A^{[p]} = A_{\aleph_0}^{[p]} \oplus A_{< \aleph_0}^{[p]}$ with

$$A_{\aleph_0}^{[p]} = \bigoplus_{\alpha_{p,n-1} = \aleph_0} \mathbb{Z}(p^n)^{(\alpha_{p,n-1})}$$

and

$$A_{< \aleph_0}^{[p]} = \bigoplus_{\alpha_{p,n-1} < \aleph_0} \mathbb{Z}(p^n)^{(\alpha_{p,n-1})},$$

we see that Corollary 5.15 applies to $A_{\aleph_0}^{[p]}$ and Proposition 5.17 to $A_{< \aleph_0}^{[p]}$; thus

$$\text{vc}^{T^{[p]}}(m) \leq m \cdot (d(U_{\geq \aleph_0}(p; A)) + 1) \quad \text{for each } m. \quad (5.7)$$

Hence from (5.5)–(5.7) and

$$A = A^{[\mathcal{P}]} \oplus \bigoplus_{p \notin \mathcal{P}} A^{[p]}$$

we see that for all m we have

$$\text{vc}^T(m) \leq ((|\mathcal{P}^c| + 1) + d) \cdot m \quad \text{where } d = \sum_p d(U_{\geq \aleph_0}(p; A)).$$

Here and below \mathcal{P}^c denotes the complement of \mathcal{P} in the set of prime numbers; i.e., \mathcal{P}^c is the (finite) set of primes singular for A . Thus (3) holds (with $(|\mathcal{P}^c| + 1) + d$ in place of d). This finishes the proof of $(4) \Rightarrow (3)$, and hence of Theorem 5.23. \square

Implicit in the proof of this theorem are upper and lower bounds on vc^T :

Corollary 5.24. *Suppose $\text{vc}^T(1) < \infty$. Then the set \mathcal{P}^c of primes singular for A is finite, $U_{\geq \aleph_0}(p; A)$ is finite for all p , and $U_{\geq \aleph_0}(p; A) = \emptyset$ for all but finitely many p . For every m we have*

$$\max\{d, |\mathcal{P}^c|\} \cdot m \leq \text{vc}^T(m) \leq (d + (|\mathcal{P}^c| + 1)) \cdot m \quad \text{where } d = \sum_p d(U_{\geq \aleph_0}(p; A)).$$

Proof Again we can assume that A is strict Szmielew as before. In the proof of Theorem 5.23 we already observed that \mathcal{P}^c is finite and $U_{\geq \aleph_0}(p; A)$ is finite for all p and non-empty only for finitely many p , and we deduced the upper bound on vc^T . For each prime p set

$$A^{\langle p \rangle} := \bigoplus_{n > 0} \mathbb{Z}(p^n)^{(\alpha_{p,n}-1)},$$

and let

$$A' := \bigoplus_p A^{\langle p \rangle}[p^{m_p+1}] \quad \text{where } m_p = \max U_{\geq \aleph_0}(p; A) \text{ with } \max \emptyset = -1.$$

Then $U_{\geq \aleph_0}(p; A') = U_{\geq \aleph_0}(p; A)$ for each p , and A' is a p.p. definable subgroup of the pure subgroup $\bigoplus_p A^{\langle p \rangle}$ of A , hence $\text{vc}^T \geq \text{vc}^{T'}$ where $T' = \text{Th}(A')$. Corollary 5.15 thus yields $\text{vc}^T(m) \geq dm$ for every m . Similarly, set

$$A'' := \bigoplus_{p \in \mathcal{P}^c} A^{\langle p \rangle}.$$

Then A'' is a pure subgroup of A , and hence Lemma 5.22 applied to A'' shows that $\text{vc}^T(m) \geq |\mathcal{P}^c|m$ for each m . \square

Note that there are totally transcendental abelian groups with infinite VC density, e.g., $A = \bigoplus_p \mathbb{Z}(p^\infty)^{(\aleph_0)}$. Theorem 5.23 allows us to give a simple characterization of superstable abelian groups with finite VC density:

Corollary 5.25. *If T is superstable, then $\text{vc}^T(1) < \infty$ iff there are only finitely many primes p such that $A[p]$ is infinite.*

Proof We may assume again that A is strict Szmielew. Now T is superstable iff there are only finitely many pairs (p, n) such that $\sum_{i \geq n} \alpha_{p,i} + \gamma_p = \aleph_0$ (see [17, Theorem A.2.13]). In particular, if T is superstable then there are only finitely many p such that A/pA is infinite, and for each p there are only finitely many n such that $\alpha_{p,n} \neq 0$. Thus the claim follows from Theorem 5.23. \square

It would be interesting to know the precise values of the VC density function of a given abelian group satisfying one of the equivalent conditions in Theorem 5.23. This was accomplished earlier in this paper in some special cases, see, e.g., Corollary 5.15. We finish this subsection with treating another simple special case:

Corollary 5.26. *Suppose T is \aleph_1 -categorical. Then $\text{vc}^T(m) = m$ for each m . (In particular, T is dp-minimal.)*

Proof By a theorem of Macintyre (see [26] or [17, Theorem A.2.12]), $A = B \oplus C$ where B is finite and either C is divisible and $C[p]$ is finite for each prime p ; or $C = \mathbb{Z}(p^n)^{(\alpha)}$ for some prime p , some $n > 0$, and some (infinite) cardinal α . In the first case, A is non-singular and hence of U-rank 1, and the claim follows from

Corollary 4.24. In the second case, C is p.p. uniserial, hence has the VC 1 property, by Corollary 4.7. So $\text{vc}^T(m) = m$ for each m , by Lemma 4.26. \square

5.4 Dp-minimal abelian groups. As in the previous subsection we assume here that A is infinite. For $m = 1$ we can improve on Corollary 5.24:

Proposition 5.27. *T is dp-minimal iff A is elementarily equivalent to one of the following abelian groups:*

1. *a direct sum of a non-singular abelian group with a group of the form*

$$(\bigoplus_{n>0} \mathbb{Z}(p^n)^{(\alpha_{n-1})}) \oplus \mathbb{Z}(p^\infty)^{(\beta)} \oplus \mathbb{Z}_{(p)}^{(\gamma)},$$

for some prime p and cardinals $\alpha_{n-1}, \beta, \gamma$, with each α_{n-1} finite;

2. *$\mathbb{Z}(p^k)^{(\alpha)} \oplus \mathbb{Z}(p^{k+1})^{(\beta)} \oplus B$ where p is a prime, $k > 0$, α, β are cardinals, at least one of which is infinite, and B is a finite abelian group.*

Note that by Corollary 5.16, the groups in (2) are precisely the \aleph_0 -categorical dp-minimal abelian groups.

We first show that each of the groups described in (1) is dp-minimal. Recall the notation introduced in the proof of Lemma 5.18: for $d < d'$, $\delta_{d,d'}^p(x)$ denotes the p.p. formula $p^{d'}|p^d x$, and $\tau_d^p(x)$ the p.p. formula $p^d x = 0$. (If p is understood from the context we drop the superscript p in this notation.) Note that if B is a non-singular abelian group and $\varphi(x)$ is a p.p. formula of the form $ax = 0$ where $a \in \mathbb{Z}$, $a \neq 0$, then $\varphi(B) \sim 0$, whereas if $\varphi = \delta_{d,d'}^p$ for some prime p and $d < d'$, then $\varphi(B) \sim B$.

Lemma 5.28. *Suppose $A = B \oplus C$ where B is non-singular and*

$$C = \bigoplus_{n>0} \mathbb{Z}(p^n)^{(\alpha_{n-1})} \oplus \mathbb{Z}(p^\infty)^{(\beta)} \oplus \mathbb{Z}_{(p)}^{(\gamma)},$$

where p is a prime and $\alpha_{n-1}, \beta, \gamma$ are cardinals, with each α_{n-1} finite. Then $\widehat{\text{PP}}(A)$ is a chain.

Proof By the remark preceding this lemma it is enough to show that

$$\tau_e^p(C) \lesssim \delta_{d,d'}^p(C) \quad \text{for all } d < d' \text{ and all } e.$$

Since

$$\delta_{d,d'}^p(\mathbb{Z}(p^\infty)) = \mathbb{Z}(p^\infty), \quad \delta_{d,d'}^p(\mathbb{Z}_{(p)}) = p^{d'-d} \mathbb{Z}_{(p)}$$

as well as

$$\tau_e^p(\mathbb{Z}(p^\infty)) = \mathbb{Z}(p^\infty)[p^e], \quad \tau_e^p(\mathbb{Z}_{(p)}) = 0,$$

this follows from (the proof of) Lemma 5.18. \square

Before we show that conversely, if T is dp-minimal, then it is the complete theory of one of the groups described in (1) and (2) of Proposition 5.27, we note:

Lemma 5.29. *Let β, γ be cardinals and $k > 0$ be an integer, and suppose*

$$A = \mathbb{Z}(p^k)^{(\aleph_0)} \oplus \mathbb{Z}(p^\infty)^{(\beta)} \oplus \mathbb{Z}_{(p)}^{(\gamma)} \oplus B$$

where B is finite. Then A is dp-minimal iff $\beta = \gamma = 0$.

Proof Since A is dp-minimal iff A/B is dp-minimal, by parts (4) and (5) of Corollary 4.19, we may assume that $B = 0$. So if $\beta = \gamma = 0$, then $A = \mathbb{Z}(p^k)^{(\aleph_0)}$ is p.p. uniserial and hence dp-minimal (by Corollaries 4.7 and 4.12). Conversely, suppose A is dp-minimal. We have

$$A[p^{k-1}] = p\mathbb{Z}(p^k)^{(\aleph_0)} \oplus (\mathbb{Z}(p^\infty)[p^{k-1}])^{(\beta)}$$

and

$$p^k A = \mathbb{Z}(p^\infty)^{(\beta)} \oplus p^k \mathbb{Z}_{(p)}^{(\gamma)}$$

with intersection

$$A[p^{k-1}] \cap p^k A = (\mathbb{Z}(p^\infty)[p^{k-1}])^{(\beta)},$$

so

$$A[p^{k-1}]/(A[p^{k-1}] \cap p^k A) \cong p\mathbb{Z}(p^k)^{(\aleph_0)}$$

and

$$p^k A/(A[p^{k-1}] \cap p^k A) \cong (\mathbb{Z}(p^\infty)/\mathbb{Z}(p^\infty)[p^{k-1}])^{(\beta)} \oplus p^k \mathbb{Z}_{(p)}^{(\gamma)}.$$

Thus $A[p^{k-1}] \cap p^k A$ has infinite index in both $A[p^{k-1}]$ and $p^k A$ iff $k > 1$ and β or γ are non-zero. By the equivalence of (1) and (4) in Corollary 4.12, this shows that, as A is dp-minimal, we have $k = 1$ or $\beta = \gamma = 0$. Suppose $k = 1$; then

$$A[p] = \mathbb{Z}(p)^{(\aleph_0)} \oplus \mathbb{Z}(p^\infty)^{(\beta)}[p], \quad pA = \mathbb{Z}(p^\infty)^{(\beta)} \oplus p\mathbb{Z}_{(p)}^{(\gamma)}, \quad A[p] \cap pA = \mathbb{Z}(p^\infty)^{(\beta)}[p],$$

so $A[p] \cap pA$ has infinite index in $A[p]$, and $A[p] \cap pA$ has infinite index in pA iff $\beta > 0$ or $\gamma > 0$. Again, dp-minimality of A and Corollary 4.12 yield $\beta = \gamma = 0$. \square

Now suppose T is dp-minimal; so in particular $\text{vc}^T(1) = 1$ by Corollary 4.12. By Corollary 5.24, this implies that there is at most one prime singular for A , and $d(U_{\geq \aleph_0}(p; A)) \leq 1$ for all p , with $d(U_{\geq \aleph_0}(p; A)) = 1$ for at most one p . (Recall the definition of $d(\cdot)$ from Section 5.1.) If no prime is singular for A , then A is of type (1) (with p arbitrary and $\alpha_{n-1} = \beta = \gamma = 0$ for all $n > 0$), so we may assume that p is the unique prime which is singular for A .

We may and shall assume that A is strict Szmielew as described before the proof of Theorem 5.23. Employing the notation introduced in that proof, we then have

$$A = A^{[\mathcal{P}]} \oplus A_{\aleph_0}^{[p]} \oplus A_{< \aleph_0}^{[p]} \oplus \mathbb{Z}(p^\infty)^{(\beta_p)} \oplus \mathbb{Z}_{(p)}^{(\gamma_p)}$$

where $A^{[\mathcal{P}]}$ is non-singular. Note that by Corollary 4.19, (2), each of the direct summands of A is also dp-minimal. If $U_{\geq \aleph_0}(p; A) = \emptyset$, then $A_{\aleph_0}^{[p]} = 0$, and A is as described in (1). So from now on suppose that $U_{\geq \aleph_0}(p; A)$ is non-empty, hence $d(U_{\geq \aleph_0}(p; A)) = 1$, that is, $U_{\geq \aleph_0}(p; A) = \{k\}$ or $U_{\geq \aleph_0}(p; A) = \{k, k+1\}$, for some $k > 0$. Then, applying Lemma 5.29 to the direct summand

$$\mathbb{Z}(p^k)^{(\aleph_0)} \oplus \mathbb{Z}(p^\infty)^{(\beta_p)} \oplus \mathbb{Z}_{(p)}^{(\gamma_p)}$$

of A , we see that $\beta_p = \gamma_p = 0$. Thus, to show that A is of type (2), it is enough to prove that $B := A^{[\mathcal{P}]} \oplus A_{< \aleph_0}^{[p]}$ is finite, and for this, after passing to a direct summand, we may assume that

$$A = \mathbb{Z}(p^k)^{(\aleph_0)} \oplus B.$$

Then by (the proof of) Lemma 5.18 and the remark before Lemma 5.28 we have

$$\delta_{k-1, k}^p(A) \sim p\mathbb{Z}(p^k)^{(\aleph_0)} \oplus A^{[\mathcal{P}]} \oplus pA_{< \aleph_0}^{[p]}$$

whereas

$$\tau_k^p(A) = A[p^k] \sim \mathbb{Z}(p^k)^{(\aleph_0)} \oplus (A_{<\aleph_0}^{[p]})[p^k].$$

Hence

$$\delta_{k-1,k}^p(A) \cap \tau_k^p(A) \sim p\mathbb{Z}(p^k)^{(\aleph_0)} \oplus (A_{<\aleph_0}^{[p]})[p^k]$$

has infinite index in $\tau_k^p(A)$. Now suppose for a contradiction that B is infinite. Then $\delta_{k-1,k}^p(A) \cap \tau_k^p(A)$ also has infinite index in $\delta_{k-1,k}^p(A)$, contradicting dp-minimality of A . This shows that indeed, B is finite.

This finishes the proof of Proposition 5.27. \square

Corollary 5.30. *Suppose T is totally transcendental. Then T is dp-minimal iff A is elementarily equivalent to the direct sum of a finite abelian group with one of the following:*

1. *an infinite divisible group D such that $D[p]$ infinite for at most one p ;*
2. *$\mathbb{Z}(p^k)^{(\aleph_0)}$ where p is a prime and $k > 0$;* or
3. *$\mathbb{Z}(p^k)^{(\aleph_0)} \oplus \mathbb{Z}(p^{k+1})^{(\aleph_0)}$ where p is a prime and $k > 0$.*

Proof By a theorem of Macintyre (see [26] or [17, Theorem A.2.11]), $A = B \oplus D$ where B has finite exponent and D is divisible. The claim now follows from Proposition 5.27 and the observation that every non-singular abelian group of finite exponent is finite. \square

By [28], a stable theory is dp-minimal iff all 1-types have weight 1. It would be interesting to obtain a proof of Proposition 5.27 using only the methods of geometric stability theory. It is instructive to see how Proposition 5.27 entails the well-known fact (cf. [28, Fact 3.2]) that every U-rank 1 abelian group is dp-minimal. The following lemma shows, in fact, that if a singular abelian group has U-rank 1, then it is of type (2) in Proposition 5.27.

Lemma 5.31. *A has U-rank 1 iff A is non-singular or of the form $A = \mathbb{Z}(p)^{(\aleph_0)} \oplus B$ where p is a prime and B is finite.*

Proof Clearly if A is non-singular, or if $A = \mathbb{Z}(p)^{(\aleph_0)} \oplus B$ where p is a prime and B is finite, then A has U-rank 1. So assume conversely that A has U-rank 1, and let p be a prime singular for A . We are allowed to assume that A is strict Szmielew as before. Suppose first that $A[p]$ is infinite. Then $A/A[p]$ is finite, since A has U-rank 1. This yields $\alpha_{p,n-1} < \aleph_0$ for all $n > 1$, with $\alpha_{p,n-1} = 0$ for all but finitely many $n > 1$, $\beta_p = \gamma_p = \nu = 0$, and $\bigoplus_{q \neq p} A^{[q]}$ is finite. Hence A is a direct sum of $\mathbb{Z}(p)^{(\alpha_{p,0})}$ with a finite abelian group, as required. Similarly, if A/pA is infinite, then pA is finite, and again we obtain the same conclusion. \square

5.5 Dp-minimal expansions of the group of integers. The abelian group $(\mathbb{Z}, +)$ of integers, having U-rank 1, is dp-minimal. In fact, in [1] we showed that the ordered abelian group $(\mathbb{Z}, <, +)$ of integers is also dp-minimal, and has no proper dp-minimal expansions. One might wonder whether the non-dp-minimality of the structures in question is caused by the presence of the ordering (due to our use, in the proof of this fact given in [1], of a lemma from [40], which specifically deals with ordered dp-minimal structures). We do not know the answer to the following:

Question 5.32. *Is every dp-minimal expansion of $(\mathbb{Z}, +)$ definable in $(\mathbb{Z}, <, +)$?*

We have a positive answer in several examples (which are well-known to have otherwise good model-theoretic properties). This is based on the following observation. (Below we will apply this lemma in the case $U = V$.)

Lemma 5.33. *Let $(A, +)$ be a commutative semigroup, let $U \subseteq V$ be infinite subsets of A , and $C > 0$. Suppose for each $u \in U$ there are at most C elements u' of U such that $u + u' = v + v'$ for some $v, v' \in V$ with $\{u, u'\} \neq \{v, v'\}$. Then $A = (A, +, V)$ is not dp-minimal.*

Proof Let $I = \{u, u'\}$ range over the 2-element subsets of U . Let us say that I is *good* if the only representation of $u + u'$ as a sum of two elements of V is the given one, that is, if $v, v' \in V$ with $u + u' = v + v'$, then $\{u, u'\} = \{v, v'\}$. If I is good, then $u + u' \notin v + V$, for each $v \in V \setminus I$. Consider the partitioned formula

$$\varphi(x; y) = \exists z(z \in V \wedge x = y + z)$$

in the language of A . By what we have shown, for each choice of good $I = \{u, u'\}$ the element $u + u'$ satisfies

$$\{\varphi(x; u), \varphi(x; u')\} \cup \{\neg\varphi(x; v) : v \in V \setminus I\}.$$

Each $u \in U$ is contained in no more than C 2-element subsets of U which are not good. Hence given $U' \subseteq U$ with $|U'| = n$, at least $\binom{n}{2} - Cn$ 2-element subsets of U' are good. Since for each $r > 0$ and C' we have $\binom{n}{2} - Cn > C'n^{2-1/r}$ for $n \gg 0$, we see by Lemma 2.8 that A is not dp-minimal. \square

Recall that a set U of elements of a commutative semigroup, written additively, is said to be a *Sidon set* if every element of $U + U$ can be written as a sum of two elements of U in essentially a unique way, that is: if $u + v = u' + v'$ where $u, v, u', v' \in U$, then $\{u, v\} = \{u', v'\}$. (In the terminology of the proof above, all two-element subsets of U are good.) We say that a set U is a *weak Sidon set* if for all $u, v, u', v' \in U$ with $u + v = u' + v'$ we have $u = v$ or $u' = v'$ or $\{u, v\} = \{u', v'\}$. From the previous lemma we immediately obtain:

Corollary 5.34. *Let U be an infinite weak Sidon set in a commutative semigroup $(A, +)$. Then $(A, +, U)$ is not dp-minimal.*

Example. If R is a commutative ring, M is an R -module, and U an infinite R -linearly independent subset of M , then U is a weak Sidon set (and a Sidon set if R is not of characteristic 2); hence $(M, +, U)$ is not dp-minimal. (This applies, e.g., to the structures introduced in the proof of Proposition 4.10 in [1].)

Fast growing integer sequences form Sidon sets. More precisely, let U be an infinite subset of \mathbb{Z} , enumerated as $u_0 < u_1 < \dots < u_n < u_{n+1} < \dots$, and suppose $u_0 > 0$ and $u_{n+1} \geq 2u_n$ for each n ; then U is a Sidon set.

Corollary 5.35. *The structure $(\mathbb{Z}, +, U)$ is not dp-minimal if U is one of the following:*

1. $U = \{b^n : n \geq 0\}$ or $U = \{n + b^n : n \geq 0\}$, where $b \in \mathbb{N}$, $b > 1$;
2. $U = \{n! : n \geq 0\}$; or
3. $U = \{F_n : n \geq 0\}$ where F_n is the n th Fibonacci number: $F_0 = F_1 = 1$, and $F_{n+2} = F_{n+1} + F_n$ for each n .

Proof For parts (1) and (2) use Corollary 5.34 and the fact that the set of powers b^n and the set of integers of the form $n + b^n$, where $b \in \mathbb{N}$, $b > 1$, as well as the set of factorials, are examples of Sidon sets. The set of Fibonacci numbers is not a Sidon set, but it is a weak Sidon set, and hence Corollary 5.34 also applies. In fact, using Zeckendorf's Theorem (every positive integer can be represented in a unique way as the sum of distinct Fibonacci numbers, no two of which are consecutive) one shows easily that $a \in \mathbb{N}$ has two representations as the sum of two Fibonacci numbers iff $a = 2F_n$ for some $n \geq 2$, and in this case $a = F_{n+1} + F_{n-2}$ is the only other representation of a as a sum of two Fibonacci numbers. \square

In the examples considered in the previous corollary, the corresponding expansion of the ordered abelian group $(\mathbb{Z}, <, +)$ of integers by a predicate symbol for U has quantifier elimination in a natural expansion of the language $\{<, +, U\}$ (see [8; 30]), and one can show that these expansions of $(\mathbb{Z}, <, +)$ are NIP.

We finish by noting that Lemma 5.33 also yields a result about expansions of fields by a subgroup of their multiplicative group. For this, let K be a field and G be an infinite subgroup of K^\times .

Corollary 5.36. *If G has the Mann property, then $(K, +, G)$ is not dp-minimal.*

Recall that G is said to have the *Mann property* if every equation

$$a_1x_1 + \cdots + a_nx_n = 1$$

where $n \geq 2$ and a_1, \dots, a_n are non-zero coefficients in the prime field of K , has only finitely many non-degenerate solutions in G , i.e., solutions $(g_1, \dots, g_n) \in G^n$ such that $\sum_{i \in I} a_ig_i \neq 0$ for each non-empty $I \subseteq [n]$. (Examples for multiplicative groups with the Mann property include all finite-rank subgroups of K^\times if K is algebraically closed of characteristic zero, by [12].) In [10], van den Dries and Günaydin study the model theory of pairs (K, G) (in the language of fields expanded by a unary predicate symbol) where K is algebraically closed or real closed and G has the Mann property; in particular, they show that if K is algebraically closed and G has the Mann property, then (K, G) is stable [10, Corollary 6.2].

To prove Corollary 5.36, suppose G has the Mann property and let C be a bound on the number of non-degenerate solutions of the equation $x_1 + x_2 - x_3 = 1$ in G . Then for each $u, v, u', v' \in G$ with $u + v = u' + v'$ and $\{u, v\} \neq \{u', v'\}$, either $u + v = 0$ or the triple $(u'/u, v'/u, v/u)$ is a non-degenerate solution of $x_1 + x_2 - x_3 = 1$. Hence taking $(K, +)$, G , and $C + 1$ for $(A, +)$, U and C , respectively, the hypothesis of Lemma 5.33 is satisfied. \square

Note that Corollary 5.36 is mainly interesting if K is not real closed, since Simon [40, Corollary 3.7] has shown that an expansion of a (linearly) ordered group is o-minimal if it is dp-minimal and definably complete and its underlying group is divisible. (These conditions are obviously necessary for o-minimality.) In particular, the only dp-minimal expansions of the ordered group $(\mathbb{R}, +, <)$ of real numbers are the o-minimal ones.

If K is algebraically closed and G has the Mann property, then (K, G) is not dp-minimal; this consequence of Corollary 5.36 may be strengthened as follows. (That a sufficiently saturated structure $\mathcal{M} = (K, G)$ where K is algebraically closed satisfies the hypothesis of the lemma below with $X = G$ follows from [10, Lemmas 2.2 (3) and 6.1].)

Lemma 5.37. *Let M be a structure which expands a field. Suppose that X is an infinite definable subset of M so that the transcendence degree of M over its subfield generated by X is at least n . Then there is a definable one-to-one function $X^{n+1} \rightarrow M$, hence the dp-rank of $\text{Th}(M)$ is at least $n+1$.*

Proof Pick elements $e_1, \dots, e_n \in M$ which are algebraically independent over the subfield of M generated by X . Consider the function $X^{n+1} \rightarrow M$ given by $(x_1, \dots, x_{n+1}) \mapsto x_1e_1 + \dots + x_ne_n + x_{n+1}$. It is easy to see that this function is one-to-one. The claim thus follows from Lemma 2.10. \square

We plan to systematically investigate the VC density functions of pairs of structures (including the expansions of fields by groups with the Mann property studied in [10]) at another occasion.

References

- [1] M. Aschenbrenner, A. Dolich, D. Haskell, D. Macpherson, S. Starchenko, *Vapnik-Chervonenkis density in some theories without the independence property, I*, preprint (2011). 2, 3, 5, 6, 7, 8, 9, 16, 18, 21, 22, 24, 25, 28, 41, 47, 48
- [2] P. Assouad, *Densité et dimension*, Ann. Inst. Fourier (Grenoble) **33** (1983), no. 3, 233–282. 2
- [3] J. Baldwin, α_T is finite for \aleph_1 -categorical T , Trans. Amer. Math. Soc. **181** (1973), 37–51. 14
- [4] A. Baudisch, *Magidor-Malitz quantifiers in modules*, J. Symbolic Logic **49** (1984), 1–8. 28
- [5] O. Belegradek, Y. Peterzil, F. Wagner, *Quasi-o-minimal structures*, J. Symbolic Logic **65** (2000), no. 3, 1115–1132. 15
- [6] G. Birkhoff, *Lattice Theory*, 3rd ed., American Mathematical Society Colloquium Publications, vol. XXV, American Mathematical Society, Providence, R.I., 1967. 18
- [7] G. Cherlin, L. Harrington, A. H. Lachlan, \aleph_0 -categorical, \aleph_0 -stable structures, Ann. Pure Appl. Logic **28** (1985), no. 2, 103–135. 14
- [8] G. Cherlin, F. Point, *On extensions of Presburger arithmetic*, in: B. I. Dahn (ed.), *Proceedings of the fourth Easter conference on model theory (Gross Köris, 1986)*, pp. 17–34, Seminarberichte, vol. 86, Humboldt Universität, Sektion Mathematik, Berlin, 1986. 49
- [9] A. Dolich, J. Goodrick, D. Lippel, *Dp-minimal theories: basic facts and examples*, Notre Dame J. Formal Logic **52** (2011), no. 3, 267–288. 2, 5, 9, 15
- [10] L. van den Dries, A. Günaydin, *The fields of real and complex numbers with a small multiplicative group*, Proc. London Math. Soc. (3) **93** (2006), no. 1, 43–81. 49, 50
- [11] P. Eklof, I. Herzog, *Model theory of modules over a serial ring*, Ann. Pure Appl. Logic **72** (1995), no. 2, 145–176. 21
- [12] J.-H. Evertse, H. P. Schlickewei, and W. M. Schmidt, *Linear equations in variables which lie in a multiplicative group*, Ann. of Math. (2) **155** (2002), no. 3, 807–836. 49

- [13] V. Guingona, *On uniform definability of types over finite sets*, J. Symbolic Logic, to appear. 7
- [14] G. Grätzer, *General Lattice Theory*, 2nd ed., Birkhäuser Verlag, Basel, 1998. 19
- [15] P. Grzeszczuk, E. R. Puczyłowski, *On Goldie and dual Goldie dimensions*, J. Pure Appl. Algebra **31** (1984), no. 1-3, 47–54. 18
- [16] E. Harzheim, *Ordered Sets*, Advances in Mathematics (Springer), vol. 7, Springer, New York, 2005. 17
- [17] W. Hodges, *Model Theory*, Encyclopedia of Mathematics and its Applications, vol. 42, Cambridge University Press, Cambridge, 1993. 17, 25, 41, 42, 44, 47
- [18] E. Hrushovski, K. Peterzil, A. Pillay, *Groups, measures, and the NIP*, J. Amer. Math. Soc. **21** (2008), no. 2, 563–596. 2
- [19] E. Hrushovski, A. Pillay, *Weakly normal groups*, in: Ch. Berline et al. (eds.), *Logic Colloquium '85 (Orsay, 1985)*, pp. 233–244, Stud. Logic Found. Math., vol. 122, North-Holland, Amsterdam, 1987. 23
- [20] ———, *On NIP and invariant measures*, J. Eur. Math. Soc. **13** (2011), no. 4, 1005–1061. 2
- [21] E. Hrushovski, A. Pillay, P. Simon, *Generically stable and smooth measures in NIP theories*, preprint (2010). 2
- [22] T. Kővari, V. T. Sós, and P. Turán, *On a problem of K. Zarankiewicz*, Colloquium Math. **3** (1954), 50–57. 8
- [23] D. Lascar, *Les groupes ω -stables de rang fini*, Trans. Amer. Math. Soc. **292** (1985), no. 2, 451–462. 15
- [24] M. C. Laskowski, *Vapnik-Chervonenkis classes of definable sets*, J. London Math. Soc. (2) **45** (1992), no. 2, 377–384. 2
- [25] ———, unpublished notes. 15
- [26] A. J. Macintyre, *On ω_1 -categorical theories of abelian groups*, Fund. Math. **70** (1971), 253–70. 44, 47
- [27] J. Matoušek, E. Welzl, L. Wernisch, *Discrepancy and approximations for bounded VC-dimension*, Combinatorica **13** (1993), no. 4, 455–66. 2
- [28] A. Onshuus, A. Usvyatsov, *On dp-minimality, strong dependence and weight*, J. Symbolic Logic **76** (2011), no. 3, 737–758. 2, 8, 14, 24, 47
- [29] A. Pillay, W. Y. Pong, *On Lascar rank and Morley rank of definable groups in differentially closed fields*, J. Symbolic Logic **67** (2002), no. 3, 1189–1196. 14
- [30] F. Point, *On decidable extensions of Presburger arithmetic: from A. Bertrand numeration systems to Pisot numbers*, J. Symbolic Logic **65** (2000), no. 3, 1347–1374. 49
- [31] B. Poizat, *Une preuve par la théorie de la déviation d'un théorème de J. Baldwin*, C. R. Acad. Sci. Paris Sér. A-B **287** (1978), no. 8, A589–A591. 14

- [32] M. Prest, *Rings of finite representation type and modules of finite Morley rank*, J. Algebra **88** (1984), 502–533. 21
- [33] ———, *Model Theory and Modules*, London Mathematical Society Lecture Note Series, vol. 130, Cambridge University Press, Cambridge, 1988. 17, 20, 27, 28, 41
- [34] J. Saffe, *Categoricity and ranks*, J. Symbolic Logic **49** (1984), no. 4, 1379–1392. 14
- [35] N. Sauer, *On the density of families of sets*, J. Combinatorial Theory Ser. A **13** (1972), 145–147. 2, 3
- [36] S. Shelah, *Stability, the f.c.p., and superstability; model theoretic properties of formulas in first-order theory*, Ann. Math. Logic **3** (1971), no. 3, 271–362. 2
- [37] ———, *A combinatorial problem; stability and order for models and theories in infinitary languages*, Pacific J. Math. **41** (1972), 247–261. 2, 3
- [38] ———, *Classification Theory and the Number of Nonisomorphic Models*, 2nd ed., Studies in Logic and the Foundations of Mathematics, vol. 92, North-Holland Publishing Co., Amsterdam, 1990. 4, 10, 11, 13, 14, 27
- [39] ———, *Strongly dependent theories*, Israel J. Math., to appear, available online at <http://arxiv.org/abs/math/0504197>. 2, 8
- [40] P. Simon, *On dp-minimal ordered structures*, J. Symbolic Logic **76** (2011), no. 2, 448–460. 47, 49
- [41] W. T. Trotter, *Combinatorics and Partially Ordered Sets. Dimension Theory*, Johns Hopkins Series in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, 1992. 17, 18, 19
- [42] V. N. Vapnik, A. Ja. Červonenkis, *The uniform convergence of frequencies of the appearance of events to their probabilities*, Theor. Probability Appl. **16** (1971), 264–280. 2, 3

Acknowledgments

Work on this paper was begun while some of the authors were participating in the thematic program on O-minimal Structures and Real Analytic Geometry at the Fields Institute in Toronto (Spring 2009), and attending the Durham Symposium on New Directions in the Model Theory of Fields (July 2009), organized by the London Mathematical Society and funded by EPSRC grant EP/F068751/1. The support of these institutions is gratefully acknowledged. Aschenbrenner was partly supported by NSF grant DMS-0556197. He would also like to express his gratitude to Andreas Baudisch and Humboldt-Universität Berlin for their hospitality during Fall 2010. Haskell’s research was supported by NSERC grant 238875, Macpherson acknowledges support by EPSRC grant EP/F009712/1, and Starchenko was partly supported by NSF grant DMS-0701364. We also thank the referees for the careful reading of the manuscript, in particular, for detecting an error in the original formulation and proof of Proposition 5.27.

Box 951555
Los Angeles, CA 90095-1555, U.S.A.
matthias@math.ucla.edu

Department of Mathematics and Computer Science
Kingsborough Community College
2001 Oriental Blvd.
Brooklyn NY 11235, U.S.A.
alfredo.dolich@kbcc.cuny.edu

Department of Mathematics and Statistics
McMaster University
1280 Main St W
Hamilton ON L8S 4K1, Canada
haskell@math.mcmaster.ca

School of Mathematics
University of Leeds
Leeds LS2 9JT, U.K.
h.d.macpherson@leeds.ac.uk

Department of Mathematics
University of Notre Dame
255 Hurley Building
Notre Dame, IN 46556-4618, U.S.A.
starchenko.1@nd.edu